

Annals of Computer Science and Information Systems
Volume 39

**Proceedings of the 19th Conference on
Computer Science and Intelligence Systems
(FedCSIS)**

September 8–11, 2024. Belgrade, Serbia



**Marek Bolanowski, Maria Ganzha, Leszek Maciaszek,
Marcin Paprzycki, Dominik Ślęzak (eds.)**



Annals of Computer Science and Information Systems, Volume 39

Series editors:

Maria Ganzha (Editor-in-Chief),

*Systems Research Institute Polish Academy of Sciences and Warsaw University of
Technology, Poland*

Leszek Maciaszek,

Macquarie University, Australia and Wrocław University of Economy, Poland

Marcin Paprzycki,

Systems Research Institute Polish Academy of Sciences and Management Academy, Poland

Dominik Ślęzak,

University of Warsaw, Poland and QED Software, Poland, and DeepSeas, USA

Senior Editorial Board:

Wil van der Aalst,

RWTH Aachen University, Netherlands

Enrique Alba,

University of Málaga, Spain

Marco Aiello,

University of Stuttgart, Germany

Mohammed Atiquzzaman,

University of Oklahoma, USA

Christian Blum,

Artificial Intelligence Research Institute (IIIA-CSIC), Spain

Jan Bosch,

Chalmers University of Technology, Sweden

George Boustras,

European University Cyprus, Cyprus

Barrett Bryant,

University of North Texas, USA

Rajkumar Buyya,

University of Melbourne, Australia

Chris Cornelis,

Ghent University, Belgium

Hristo Djidjev,

Los Alamos National Laboratory, USA and Bulgarian Academy of Sciences, Bulgaria

Włodzisław Duch,

Nicolaus Copernicus University, Toruń, Poland

Hans-George Fill,

University of Fribourg, Switzerland

Ana Fred,

University of Lisbon, Portugal

Giancarlo Guizzardi,

University of Twente, Netherlands

Francisco Herrera,

University of Granada, Spain

Mike Hinchey,

University of Limerick, Ireland

Janusz Kacprzyk,

Systems Research Institute, Polish Academy of Sciences, Poland

Irwin King,
The Chinese University of Hong Kong, China

Michael Luck,
King's College London, United Kingdom

Ivan Luković,
University of Belgrade, Serbia

Marjan Mernik,
University of Maribor, Slovenia

Michael Segal,
Ben-Gurion University of the Negev, Israel

Andrzej Skowron,
University of Warsaw, Poland

John F. Sowa,
VivoMind Research, LLC, USA

George Spanoudakis,
University of London, United Kingdom

Editorial Associates:

Katarzyna Wasielewska,
Systems Research Institute Polish Academy of Sciences, Poland

Paweł Sitek,
Kielce University of Technology, Poland

TeXnical editor: Aleksander Denisiuk,
University of Warmia and Mazury in Olsztyn, Poland

Promotion and Marketing: Anastasiya Danilenka,
Warsaw University of Technology, Poland

Proceedings of the 19th Conference on Computer Science and Intelligence Systems (FedCSIS)

Marek Bolanowski, Maria Ganzha, Leszek
Maciaszek, Marcin Paprzycki, Dominik Ślęzak (eds.)



2024, Warszawa,
Polskie Towarzystwo
Informatyczne



2024, New York City,
Institute of Electrical and
Electronics Engineers

Annals of Computer Science and Information Systems, Volume 39

Proceedings of the 19th Conference on Computer Science and
Intelligence Systems (FedCSIS)

ART: ISBN 978-83-969601-8-4, IEEE Catalog Number CFP2485N-ART

USB: ISBN 978-83-969601-7-7, IEEE Catalog Number CFP2485N-USB

Web: ISBN 978-83-969601-6-0

ISSN 2300-5963

DOI 10.15439/978-83-969601-6-0

© 2024, Polskie Towarzystwo Informatyczne

Ul. Solec 38/103

00-394 Warsaw, Poland

Contact: secretariat@fedcsis.org

<http://annals-csis.org/>

Cover art: Margarita

Alicja Król,

Elbląg, Poland (<https://lukawsztuce.pl>)

Also in this series:

Volume 41: Communication Papers of the 19th Conference on Computer Science and
Intelligence Systems (FedCSIS), **ISBN WEB: 978-83-969601-1-5, ISBN USB: 978-83-969601-2-2**

Volume 40: Position Papers of the 19th Conference on Computer Science and
Intelligence Systems (FedCSIS), **ISBN WEB: 978-83-969601-9-1, ISBN USB: 978-83-969601-0-8**

Volume 38: Proceedings of the Eighth International Conference on Research in
Intelligent Computing in Engineering, **ISBN WEB: 978-83-969601-5-3**

Volume 37: Communication Papers of the 18th Conference on Computer Science and
Intelligence Systems, **ISBN WEB: 978-83-969601-3-9, ISBN USB: 978-83-969601-4-6**

Volume 36: Position Papers of the 18th Conference on Computer Science and
Intelligence Systems, **ISBN WEB: 978-83-969601-1-5, ISBN USB: 978-83-969601-2-2**

Volume 35: Proceedings of the 18th Conference on Computer Science and Intelligence
Systems, **ISBN WEB 978-83-967447-8-4, ISBN USB 978-83-967447-9-1, ISBN ART 978-83-969601-0-8**

Volume 34: Proceedings of the Third International Conference on Research in
Management and Technovation **ISBN 978-83-965897-8-1**

Volume 33: Proceedings of the Seventh International Conference on Research in
Intelligent and Computing in Engineering, **ISBN WEB: 978-83-965897-6-7,**

ISBN USB: 978-83-965897-7-4

Volume 32: Communication Papers of the 17th Conference on Computer Science and
Intelligence Systems, **ISBN WEB: 978-83-965897-4-3, ISBN USB: 978-83-965897-5-0**

Volume 31: Position Papers of the 17th Conference on Computer Science and
Intelligence Systems, **ISBN WEB: 978-83-965897-2-9, ISBN USB: 978-83-965897-3-6**

Volume 30: Proceedings of the 17th Conference on Computer Science and Intelligence
Systems, **ISBN WEB: 978-83-962423-9-6, ISBN USB: 978-83-965897-0-5**

Volume 29: Recent Advances in Business Analytics. Selected papers of the 2021
KNOWCON-NSAIS workshop on Business Analytics **ISBN WEB: 978-83-962423-7-2,**

ISBN USB: 978-83-962423-6-5

DEAR Reader, it is our pleasure to present to you Proceedings of the 19th Conference on Computer Science and Intelligence Systems (FedCSIS 2024), which took place on September 8-11, 2024, in Belgrade, Serbia.

FedCSIS 2024 was chaired by Ivan Lukovic, while Dragana Makajić-Nikolić was the Chair of the Organizing Committee. This year, FedCSIS was organized by the Polish Information Processing Society (Mazovia Chapter), IEEE Poland Section Computer Society Chapter, Systems Research Institute of Polish Academy of Sciences, The Faculty of Mathematics and Information Science Warsaw University of Technology, The Faculty of Electrical and Computer Engineering of the Rzeszów University of Technology, and The Faculty of Organizational Science of the University of Belgrade.

FedCSIS 2024 was technically co-sponsored by IEEE Poland Section, IEEE Serbia and Montenegro Section, Poland Section of IEEE Computer Society Chapter, Czechoslovakia Section of IEEE Computer Society Chapter, Serbia and Montenegro Section of IEEE Computer Society Chapter, Poland Section of IEEE Systems, Man, and Cybernetics Society Chapter, Poland Section of IEEE Computational Intelligence Society Chapter, Serbia and Montenegro Section of IEEE Computational Intelligence Society Chapter, Serbia and Montenegro Section of IEEE Education Society Chapter, Serbia and Montenegro Section of IEEE Young Professionals Affinity Group, Committee of Computer Science of Polish Academy of Sciences, Informatics Association of Serbia, and Mazovia Cluster ICT.

FedCSIS 2024 was organized in collaboration with the Strategic Partner: QED Software, and sponsored by the Ministry of Science, Technological Development and Innovation, Republic of Serbia, Banca Intesa, Nelt Group, Netconomy, Elsevier, Journal of Computer Languages, ONLYOFFICE Ascensio Systems d.o.o., Beograd, MDPI and Yettel Bank.

This year, we continued adjusting the structure of the conference. Starting from 2024, FedCSIS conferences have a single Main Track with 5 Topical Areas, Thematic Sessions and, possibly, Competitions. The slightly adjusted structure emphasizes the integrity of the conference and its closeness to the issues that are crucial for the world today. Here, we recognize the fact that, over the last few years, rapid progress of various forms of computational intelligence could have been observed. As the result, broadly understood, *intelligence* which was a separate research area became part of other areas that, previously, were explored independently. As a matter of fact, today (in 2024) it is difficult to envision research (and its applications) without an intelligence component. Reflecting this, all five Topical Areas, established within the FedCSIS Main Track, while being situated within a general domain of Computer Science, represent various aspects of Intelligence Systems. Moreover, the Thematic Sessions provide focal insights into selected areas Intelligence Systems, approached from different perspectives. Even the Data Mining Competition, having strong roots in artificial intelligence, data science and machine learning, can be seen as a path toward introducing more intelligence into real-world anchored computer systems. This vision has been depicted in Figure 1.




Figure 1. FedCSIS Conference structure; birds-eye view; TS_n denotes Technical Session.

In this context, these Proceedings consist of six parts. Part 1 contains Invited Contributions. Part 2 collects Main Track full contributions (arranged alphabetically, according to the last name of the first author, with the Topical Area represented in the metadata). Part 3 contains Main Track short contributions (again, arranged alphabetically, with the Topical Area represented in the metadata). Part 4 contains full contributions, originating from the Thematic Sessions (again, arranged alphabetically, according to the last name of the first author, with the name of the Thematic Session stated in the metadata.) Part 5 collects short papers from all Thematic Sessions. Finally, Part 6 is devoted to contributions originating from the Data Mining Competition.

Keeping this in mind, let us now introduce Keynote Speakers, the remaining Invited Contributions, and the five Topical Areas of the FedCSIS 2024 Main Track.

I. INVITED CONTRIBUTIONS

FedCSIS 2024 invited four keynote speakers to deliver lectures providing a broader context for the conference participants. Moreover, two past FedCSIS keynote speakers have been invited to prepare contributions, which refer to the core focus of the conference series. There are also two contributions originating from special invited talks. The following Keynote and Invited presentations have been delivered:

- + Frank, Ulrich, University of Duisburg-Essen, Germany, *keynote title*: Multi-Level Language Architectures: Fostering Reuse, Integration and User Empowerment by Allowing for Additional Abstraction
- + Jovanović, Jelena, University of Belgrade, Serbia, *keynote title*: Learning analytics: Challenges and opportunities opened by AI
- + Kutyniok, Gitta, Ludwig-Maximilians-Universität München, Germany, *keynote title*: Reliable AI: Successes, Challenges, and Limitations
- + Tolvanen, Juha-Pekka, Metacase, Finland, *keynote title*: Languages for non-developers: what, how, where?
- + Dujmović, Jozo, San Francisco State University, USA, *invited presentation title*: Graded Logic and Professional Decision Making

II. ADVANCED ARTIFICIAL INTELLIGENCE IN APPLICATIONS

This Topical Area covers a wide range of core aspects of AI. Nowadays, AI is usually perceived as closely related to the data, therefore, the scope of this Topical Area includes, among others, elements of machine learning, data science, and big data processing, with important emerging aspects such as interactive learning and human-centered AI, as well as interpretable learning, explainable AI, and trustworthiness. Furthermore, since the realm of AI is far richer, the ultimate goal of this Topical Area is to show relationships between all of, currently pursued AI subareas, emphasizing a cross-disciplinary nature of various research branches. In 2024, the collection of papers accepted to this Topical Area has clearly reflected this cross-disciplinary nature particularly well. We can see here various areas of AI (also outside so-called “core AI”), as well as a mix of theoretical and practical contributions. From the perspective of the general scope of FedCSIS, this Topical Area embraces AI methods and examples of their applications in different practical fields.

This Topical Area was curated by:

- + Corizzo, Roberto, American University, USA
- + Sosnowski, Łukasz, Systems Research Institute of Polish Academy of Sciences, Poland
- + Szczuka, Marcin, University of Warsaw, Poland
- + Zdravevski, Eftim, Ss. Cyril and Methodius University, Macedonia

III. COMPUTER SCIENCE & SYSTEMS

This Topical Area aims at integrating and creating synergy between Computer Science and related disciplines, with Intelligence being of the core interest. The area’s scope spans themes ranging from hardware issues close to computer engineering via software issues tackled by the theory and applications of Computer Science. When compared to the previously discussed Topical Area on “Advanced Artificial Intelligence in Applications”, herein we are interested more in software system realizations and computational aspects. Therefore, we take a step from AI regarded as the set of methods, towards Intelligence Systems, understood as software systems with the elements of AI.

This Topical Area was curated by:

- + Casalino, Gabriella, University of Bari "Aldo Moro", Italy
- + Ducange, Pietro, University of Pisa, Italy
- + Pawłowski, Wiesław, University of Gdańsk, Poland
- + Wasielewska-Michniewska, Katarzyna, Systems Research Institute of Polish Academy of Sciences, Poland

IV. NETWORK SYSTEMS & APPLICATIONS

Modern network systems encompass a wide range of solutions and technologies, including wireless and wired networks, network systems, services, and applications. On the one hand, network technologies are used in the majority of areas that make human life easier and more comfortable. On the other hand, the rapid

need for network deployment brings new challenges in network management and network design, which are reflected in hardware, software, services, and security-related problems. Going back to the main scope of FedCSIS, it is obvious that network solutions are one of the crucial layers of scalable modern software systems, including Intelligence Systems. On the other hand, equally obviously, AI methods can be useful to make network systems and their applications more efficient. Accordingly, the aim of this Topical Area is to bring more Intelligence into all aspects of network systems. Moreover, besides network systems, one should think also about network models, network algorithms, etc. Therefore, this Topical Area covers not only the technological side, but also the societal and social impacts of network developments.

This Topical Area was curated by:

- + Armando, Alessandro, University of Genova, Italy
- + Awad, Ali Ismail, United Arab Emirates University, United Arab Emirates
- + Furtak, Janusz, Military University of Technology, Poland
- + Hodoň, Michal, University of Žilina, Slovakia
- + Suri, Niranjan, Institute of Human and Machine Cognition, United States

V. INFORMATION TECHNOLOGY FOR BUSINESS & SOCIETY

The aim of this Topical Area is to integrate and create synergy between disciplines of information technology (IT), Intelligence Systems, and social sciences. Collected contributions address issues relevant to IT and necessary for practical, everyday needs of business, other organizations, and society at large. Moreover, they take a socio-technical view on Intelligence Systems and, at the same time, relate to ethical, social and political issues that they raise. Thus, from the viewpoint of the FedCSIS as a whole, this Topical Area goes beyond Computer Science itself. It refers to the fact that every software system or solution, and especially a system or solution with some flavors of Intelligence, needs to be carefully deployed in real life. In other words, it is not only about machines – it is also about humans. Accordingly, this Topical Area embraces research on methods and processes of adoption of AI and Intelligence Systems in society and specific markets of business applications.

This Topical Area was curated by:

- + Cano, Alberto, Virginia Commonwealth University, Richmond, United States
- + Dias, Gonçalo, University of Aveiro, Portugal
- + Miller, Gloria, Maxmetrics, Germany
- + Naldi, Maurizio, LUMSA University, Italy
- + Wątróbski, Jarosław, University of Szczecin, Poland
- + Ziemia, Ewa, University of Economics in Katowice, Poland

VI. SOFTWARE, SYSTEM & SERVICE ENGINEERING

For decades, an open question in the software industry remains, how to provide fast and effective software process and software services, and how realize software systems,

embedded systems, autonomous systems, or cyber-physical systems that will address the open issue of supporting information management process in various, particularly complex organization systems. Even more, it is a hot issue how to provide a synergy between systems in common, and software services as a mandatory component of each modern organization, particularly in terms of IoT, Big Data, and Industry 4.0 paradigms. Therefore, the main goal of this Topical Area is to address open questions and real potentials for various applications of modern approaches and technologies to develop and implement effective software services in a support of information management and system engineering. One can see, here, a clear linkage to AI and Intelligence Systems as well.

This Topical Area was curated by:

- + Kolukısa Tarhan, Ayça, Hacettepe University, Turkey
- + Mernik, Marjan, University of Maribor, Slovenia
- + Popović, Aleksandar, University of Montenegro, Podgorica, Montenegro

VII. DATA MINING COMPETITION

FedCSIS 2024 Data Mining Competition was devoted to: Predicting Stock Trends and was the 10th data science challenge, organized within the scope of the FedCSIS conference series. In this anniversary edition, the task was related to financial data - participants were asked to predict the performance of investments in selected stocks, from several industry sectors. The competition was sponsored by Yettel.Bank and the FedCSIS conference.

The 2024 competition was organized by:

- + Aleksandar M. Rakićević, University of Belgrade
- + Pavle D. Milošević, University of Belgrade
- + Ivana T. Dragović, University of Belgrade
- + Ana M. Poledica, University of Belgrade
- + Milica M. Zukanović, University of Belgrade
- + Ivan S. Luković, University of Belgrade
- + Andrzej Janusz, Queensland University of Technology and QED Software
- + Dominik Ślęzak, QED Software and University of Warsaw

This year, 194 teams comprising of 259 individuals registered for the competition, making it one of the most popular competitions in the history of the FedCSIS conference series. By the end of the competition, 77 enrolled teams were deemed active. Their members represented 28 different countries from around the world, with the highest representation from Germany (58), Poland (50), Italy (41), Turkey (24), and Serbia (18). There were 3,000 submitted solutions in total. Out of these, on the final day of the competition, over 250 solutions have been submitted. After evaluation, the following contributions, found in these proceedings, discuss the winning contributions:

First place: Carlos Huertas, Gradient Boosting Trees and Large Language Models for Tabular Data Few-Shot Learning

Second place: Quang Hieu Vu, Dymitr Ruta, Ling Cen and Ming Liu, FedCSIS 2024 Data Science Challenge: Predicting Stock Trends by a Multi-Dimensional Approach

Third place: Chang Lin, Key Financial Indicators Analysis and Stock Trend Forecasting Based on a Wrapper Feature Selection Method

Special award for the most practically applicable solution: Marcin Traskowski and Eyad Kannout, Forecasting Stock Trends with Feedforward Neural Networks

VIII. ZDZISŁAW PAWLAK AWARDS

The above-described five Topical Areas of the FedCSIS Main Track reflect five fundamental aspects of understanding, developing, and applying Intelligence Systems. This topical integrity is emphasized by the Professor Zdzisław Pawlak award, considered in four categories: Best Paper, Young Researcher, Industry Cooperation, and International Cooperation. Although Professor Zdzisław Pawlak has been often recognized as “the father of Polish AI”, his research achievements have gone far beyond AI itself, in particular toward AI applications and Intelligence Systems as we mean them. Accordingly, for this award contributions from the Main Track and from all Thematic Sessions are considered.

This year, the following contributions have been awarded:

- In the category **Best Paper:** Rytis Maskeliunas and Robertas Damasevicius, “d'Alembert Convolution for Enhanced Spatio-Temporal Analysis of Forest Ecosystems”
- In the category **Young Researcher:** Alexander Kammerer, Florian Burger, Daniel Lübbert and Katinka Wolter, “HPC operation with time-dependent cluster-wide power capping”
- In the category **Industry Cooperation:** Guillaume Hutzler, Hanna Klaudel, Witold Klaudel, Franck Pommereau and Artur Rataj, “An autonomous vehicle in a connected environment: case study of cyber-resilience”
- In the category **International Cooperation Award:** Alexander Fischer, Juha-Pekka Tolvanen and Ramin Tavakoli Kolagari, “Automotive Cybersecurity Engineering with Modeling Support”

Young Researcher Award and International Cooperation Award were sponsored by MDPI, while the remaining awards were sponsored by Mazovia Branch of Polish Information Processing Society.

IX. STATISTICS

Each contribution, found in this volume, was refereed by at least two referees and the acceptance rate of regular full papers was approximately 20% (37 accepted contributions, out of 184 submissions). Here, the long-term trend is depicted in Figure 2.




Figure 2. Acceptance rate for the regular full papers for the FedCSIS conference series since 2005 (when FedCSIS' predecessor was organized for the first time).

X. COMMITTEES

The Senior Program Committee of FedCSIS 2024 consisted of:

- van der Aalst, Wil, RWTH Aachen University, Germany
- Alba, Enrique, University of Málaga, Spain
- Aiello, Marco, University of Stuttgart, Germany
- Armando, Alessandro, University of Genova, Italy
- Atiquzzaman, Mohammed, University of Oklahoma, USA
- Awad, Ali Ismail, United Arab Emirates University, United Arab Emirates
- Blum, Christian, Artificial Intelligence Research Institute (IIIA-CSIC), Spain
- Bosch, Jan, Chalmers University of Technology, Sweden
- Boustras, George, European University, Cyprus
- Bryant, Barrett, University of North Texas, USA
- Buyya, Rajkumar, University of Melbourne, Australia
- Cano, Alberto, Virginia Commonwealth University, United States
- Casalino, Gabriella, University of Bari "Aldo Moro", Italy
- Corizzo, Roberto, American University, USA
- Cornelis, Chris, Ghent University, Belgium
- Dias, Gonçalo, University of Aveiro, Portugal
- Djidjev, Hristo, Los Alamos National Laboratory, USA and Institute of Information and Communication Technologies, Bulgaria
- Ducange, Pietro, University of Pisa, Italy
- Duch, Włodzisław, Nicolaus Copernicus University, Poland
- Fill, Hans-George, University of Fribourg, Switzerland
- Fred, Ana, Instituto Superior Técnico (IST—Technical University of Lisbon), Portugal
- Furtak, Janusz, Military University of Technology, Poland
- Giancarlo Guizzardi, Free University of Bolzano-Bozen, Italy
- Herrera, Francisco, University of Granada, Spain
- Hinchey, Mike, Lero, University of Limerick, Ireland
- Hodoň, Michal, University of Žilina, Slovakia
- Kacprzyk, Janusz, Systems Research Institute, Polish Academy of Sciences, Poland
- King, Irwin, The Chinese University of Hong Kong, China
- Kolukısa Tarhan, Ayça, Hacettepe University, Turkey
- Komorowski, Jan, Uppsala University, Sweden
- Kwaśnicka, Halina, Wrocław University of Science and Technology, Poland
- Luck, Michael, University of Sussex, United Kingdom
- Mernik, Marjan, University of Maribor, Slovenia
- Michalewicz, Zbigniew, University of Adelaide, Australia
- Miller, Gloria, maxmetrics, Germany
- Naldi, Maurizio, LUMSA University, Italy
- Pawłowski, Wiesław, University of Gdańsk and Systems Research Institute, Polish Academy of Sciences, Poland
- Pedrycz, Witold, University of Alberta, Canada
- Popović, Aleksandar, University of Montenegro, Montenegro
- Raś, Zbigniew, University of North Carolina, United States
- Segal, Michael, Ben-Gurion University of the Negev, Israel
- Skowron, Andrzej, Systems Research Institute, Polish Academy of Sciences, Poland
- Słowiński, Roman, Poznań University of Technology, Poland
- Sosnowski, Łukasz, Systems Research Institute, Polish Academy of Sciences, Poland
- Sowa, John F., VivoMind Research, LLC, USA
- Spanoudakis George, University of London, United Kingdom
- Suri, Niranjan, Institute of Human and Machine Cognition, United States
- Szczuka, Marcin, University of Warsaw, Poland
- Wasielewska-Michniewska, Katarzyna, Systems Research Institute, Polish Academy of Sciences, Poland
- Wątróbski, Jarosław, University of Szczecin, Poland
- Zdravevski, Eftim, Ss. Cyril and Methodius University, Macedonia
- Ziemia, Ewa, University of Economics in Katowice, Poland

The FedCSIS 2024 Program Committee consisted of:

- Abbas Musarat, Quaid-i-Azam University, Pakistan
- Adam Emmanuel, Université Polytechnique Hauts-de-France, France
- Adil Muhammad, ISMAR-CNR, Italy
- Ahmad Muhammad Ovais, Karlstad University, Sweden
- Ahmad Bilal, University of Warwick, United Kingdom
- Al-Naday Mays, University of Essex, United Kingdom
- Ameer Hamza Muhammad, COMSATS University Islamabad, Pakistan
- Anastassi Zacharias, ASPETE School of Pedagogical and Technological Education, Greece
- Arabas Jaroslaw, Warsaw University of Technology, Poland
- Arruda Filho Emílio José, FUMEC, Brasil
- Arshi Oroos, UPES, India
- Atanassov Krassimir, Bulgarian Academy of Sciences, Bulgaria

- Atasever Mesut, Uşak University, Turkey
- Azad Mohammad, Jouf University, Saudi Arabia
- Bacco Manlio, Institute of Information Science and Technologies, National Research Council, Italy
- Bachan Jolanta, Adam Mickiewicz University, Poland
- Badica Costin, University of Craiova, Romania
- Badica Amelia, University of Craiova, Romania
- Badr Georges, Holy Spirit University of Kaslik, Lebanon
- Balazs Krisztian, Budapest University of Technology and Economics, Hungary
- Baldán Lozano Francisco Javier, University of Granada, Spain
- Ballas Rüdiger G., Mobile University of Technology, Germany
- Banach Richard, University of Manchester, United Kingdom
- Barata José, Universidade Nova de Lisboa, Portugal
- Barbosa Joao, INESC-TEC, Portugal
- Barisic Ankica, Université Côte d'Azur, France
- Batalla Jordi Mongay, National Institute of Technology, Poland
- Belciug Smaranda, University of Craiova, Romania
- Bellinger Colin, National Research Council of Canada, Canada
- Ben-Assuli Ofir, Ono Academic College, Israel
- Białas Andrzej, Institute of Innovative Technologies EMAG, Poland
- Bicevskis Janis, University of Latvia, Latvia
- Binnewitt Johanna, BIBB, and University of Cologne, Germany
- Bjeladinovic Srdja, University of Belgrade, Serbia
- Blachnik Marcin, Silesian University of Technology, Poland
- Blasband Darius, RainCode, Belgium
- Bodyanskiy Yevgeniy, Kharkiv National University of Radio Electronics, NURE, Ukraine
- Boffa Stefania, University of Milano-Bicocca, Italy
- Borkowski Boleslaw, Warsaw University of Life Sciences, Poland
- Braumann Ulf-Dietrich, InfAI, Germany
- Brdjanin Drazen, University of Banja Luka, Bosnia & Herzegovina
- Brezovan Marius, University of Craiova, Romania
- Bridova Ivana, University of Zilina, Slovakia
- Bronselaer Antoon, Ghent University, Belgium
- Brzoza-Zajęcka Ada, AGH University of Science and Technology, Poland
- Burczynski Tadeusz, Polish Academy of Sciences, Poland
- Byrski Aleksander, AGH University Science and Technology, Poland
- Cabri Giacomo, Università di Modena e Reggio Emilia, Italy
- Calpe Maravilla Javier, Universitat de Valencia, Spain
- Carbone Roberto, FBK, Italy
- Carchiolo Vincenza, Università di Catania, Italy
- Cardoso Henrique Lopes, University of Porto, Portugal
- Castrillon-Santana Modesto, University of Las Palmas de Gran Canaria, Spain
- Cattaneo Gianpiero, University of Milano-Bicocca, Italy
- Ceci Michelangelo, University of Bari "A. Moro", Italy
- Celikovic Milan, University of Novi Sad, Serbia
- Cen Ling, Khalifa University, United Arab Emirates
- Challenger Moharram, Antwerp University, Belgium
- Chang Lin, Beijing University of Posts and Telecommunications, China
- Charytanowicz Malgorzata, Catholic University of Lublin, Poland
- Chen Haiming, Chinese Academy of Sciences, China
- Cherukuri Aswani Kumar, VIT University, India
- Chomiak-Orsa Iwona, Wroclaw University of Economics and Business, Poland
- Christozov Dimitar, American University in Bulgaria, Bulgaria
- Chudan David, Prague University of Economics and Business, Czech Republic
- Clarke Nathan, University of Plymouth, United Kingdom
- Clausberg Philipp, WIG2 Institute, Germany
- Colantonio Sara, ISTI-CNR, Italy
- Corpetti Thomas, University of Rennes, France
- Courty Nicolas, University of Bretagne Sud, France
- Coviello Giuseppe, Politecnico di Bari, Italy
- Culibrk Dubravko, University of Novi Sad, Serbia
- Cybulski Piotr, Military University of Technology, Poland
- Czarnačka-Chrobot Beata, Warsaw School of Economics, Poland
- Damasevicius Robertas, Silesian University of Technology, Poland
- Daszczuk Wiktor, Warsaw University of Technology, Poland
- De Juana-Espinosa Susana, Universidad de Alicante, Spain
- De Marinis Pasquale, UNIBA, Italy
- De Tre Guy, Ghent University, Belgium
- Dettmer Sandra, Swansea University, United Kingdom
- Dey Lipika, Tata Consulting Services, India
- Diaw Samba, LIMBI, Senegal
- Dimitrieski Vladimir, Faculty of Technical Sciences, Serbia
- Djordjevic Vuk, Institute of Field and Vegetable Crops, Serbia

- Domanska Joanna, Institute of Theoretical and Applied Informatics, Poland
- Dragovic Ivana, University of Belgrade, Serbia
- Drezewski Rafal, AGH University of Science and Technology, Poland
- Dudycz Helena, Wroclaw University of Economics, Poland
- Dutta Arpita, National University of Singapore, Singapore
- Dutta Soma, University of Warmia and Mazury in Olsztyn, Poland
- Eisenhardt Monika, University of Economics Katowice, Poland
- El-Halim Essam H. Houssein, Minia University, Egypt
- Engelbrecht Andries, University of Stellenbosch, South Africa
- Erata Ferhat, Yale University, USA
- Erol Barkana Duygun, Yeditepe University, Turkey
- Erradi Mohammed, Mohammed-V Souissi University, ENSIAS (Ecole Nationale Supérieure d'Informatique et d'Analyse des Systèmes), Morocco
- Escalona M.J., University of Seville, Spain
- Esposito Massimo, ICAR-CNR, Italy
- Fareh Messaouda, University Blida 1, Algeria
- Farooq Ali, University of Turku, Finland
- Fechner Richard, ECU Tübingen, Germany
- Felkner Anna, NASK - Research and Academic Computer Network, Poland
- Fertilj Krešimir, University of Zagreb, Croatia
- Fialko Sergiy, Cracow University of Technology, Poland
- Filipe Vítor, INESC TEC / UTAD, Portugal
- Fissaa Tarik, Institut National des Postes et Télécommunications, Morocco
- Fonseca Jose Manuel, UNINOVA, Portugal
- Fournier-Viger Philippe, University of Moncton, Canada
- Fuchs Christoph, University of Bonn, Germany
- Fuentes Alvaro, Jeonbuk National University, South Korea
- G. Barbosa Jorge, University of Porto, Portugal
- G.-Tóth Boglárka, University of Szeged, Hungary
- Gabryelczyk Renata, University of Warsaw, Poland
- Ganea Eugen, University of Craiova, Romania
- García-Mireles Gabriel, Universidad de Sonora, Mexico
- Gawkowski Piotr, Warsaw University of Technology, Poland
- Gburzyński Paweł, University of Alberta, Canada; Vistula University, Poland
- Gepner Paweł, Warsaw University of Technology, Poland
- Geri Nitza, The Open University of Israel, Israel
- Getz Laura, Bundesinstitut für Berufsbildung, Germany
- Gheisari Mehdi, Islamic Azad University, Iran
- Giraddi Shantala, BVBCET, India
- Gjoreski Hristijan, Ss. Cyril and Methodius University in Skopje, North Macedonia
- Göknil Arda, SINTEF Digital, Norway
- Gomolińska Anna, University of Białystok, Poland
- Gózdź Marek, UMCS, Poland
- Grabara Dariusz, University of Economics in Katowice, Poland
- Grabowski Mariusz, Cracow University of Economics, Poland
- Gravvanis George, Democritus University of Thrace, Greece
- Grochla Krzysztof, Institute of Theoretical and Applied Informatics of PAS, Poland
- Grzegorowski Marek, Warsaw University, Poland
- Gücük Gian-Luca, University of Hamburg, Germany
- Haerting Ralf, Hochschule Aalen, Germany
- Hakius Bettina, BTA Wiedenes, Germany
- Halawi Leila, Embry-Riddle Aeronautical University, USA
- Hamel Oussama, University Blida 1, Algeria
- Hasso Hussein, Fraunhofer FKIE, Wachtberg, Germany
- Heidler Tobias, WIG2 Institute, Germany
- Hein Kristine, BIBB, Bonn, Germany
- Helsingius Mika, Finnish Defence Research Agency, Finland
- Henry Christopher, University of Winnipeg, Canada
- Hernes Marcin, Wroclaw University of Economics and Business, Poland
- Horváth Zoltán, Eötvös Loránd University, Hungary
- Hosobe Hiroshi, Hosei University, Japan
- Hrach Christian, InfAI, Germany
- Hu Bao-Gang, Institute of Automation, Chinese Academy of Sciences, China
- Hübenthal Tobias, University of Cologne, Germany
- Hullam Gabor, Budapest University of Technology and Economics, Hungary
- Hussain Shahid, Institute of Business Administration, Pakistan
- Hussain Asad, University of Bergamo, Italy
- Lenco Dino, IRSTEA, France
- Ignaciuk Przemyslaw, Łódź University of Technology, Poland
- Iqbal Naeem, German Research Center for Artificial Intelligence, Germany
- Islam Taminul, Southern Illinois University Carbondale, USA
- Ivanovic Mirjana, University of Novi Sad, Serbia
- Jakovljević Nikša, University of Novi Sad, Serbia
- Janicki Ryszard, McMaster University, Canada
- Janicki Artur, Warsaw University of Technology, Poland
- Jarzbowicz Aleksander, Gdansk University of Technology, Poland

- Jha Khushboo, Birla Institute of Technology, India
- John Niels, WIG2 Institute, Germany
- Johnsen Frank, Norwegian Defence Research Establishment, Norway
- Jovancevic Igor, University of Montenegro, Montenegro
- Jovanovik Milos, Ss. Cyril and Methodius University in Skopje, North Macedonia
- Kaczmarek Katarzyna, University of Strathclyde, United Kingdom
- Kaloyanova Kalinka, University of Sofia, Bulgaria
- Kanciak Krzysztof, Military University of Technology, Poland
- Kania Krzysztof, University of Economics in Katowice, Poland
- Kapczyński Adrian, Silesian University of Technology, Poland
- Karnouskos Stamatis, SAP, Germany
- Kasprzak Włodzimierz, Politechnika Warszawska, Poland
- Keir Paul, University of the West of Scotland, Scotland
- Kelner Jan, Military University of Technology, Poland
- Keswani Bright, Suresh Gyan Vihar University, Jaipur, India
- Kieran Judith, CARR Communications, Ireland
- Kisacanin Branislav, University of Novi Sad, Serbia
- Klapp Iftach, Agriculture research Organization - Volcani Institute, Israel
- Kliegr Tomas, Prague University of Economics and Business, Czech Republic
- Kluza Krzysztof, AGH University of Science and Technology, Poland
- Kobayashi Haruo, Gunma University, Japan
- Kobylnski Andrzej, Warsaw School of Economics, Poland
- Koczy Laszlo T., Budapest University of Technology and Economics, Hungary
- Kokosinski Zbigniew, Cracow University of Technology, Poland
- Kolog Emmanuel Awuni, University of Ghana, Ghana
- Kordić Slavica, Faculty of Technical Sciences, Serbia
- Kosar Tomaz, University of Maribor, Slovenia
- Kosmopoulos Dimitrios, University of Patras, Greece
- Kovatcheva Eugenia, University of Library Studies and Information Technologies, Bulgaria
- Kozak Jan, University of Economics in Katowice, Poland
- Kozielski Stanislaw, Silesian University of Technology, Poland
- Kozłowski Artur, Łukasiewicz Research Network, Poland
- Krajsic Philippe, Center for Scalable Data Analytics and Artificial Intelligence, Germany
- Krawczyk Henryk, Gdańsk University of Technology, Poland
- Krawiec Krzysztof, Poznań University of Technology, Poland
- Krdzavac Nenad, Leibniz Information Centre for Science and Technology, Germany
- Krdzavac Nenad, Leibniz Information Centre for Science and Technology, Germany
- Krüger Kai, German Federal Institute for Vocational Education and Training, Germany
- Kryvyi Serhii, Taras Shevchenko National University of Kyiv, Ukraine
- Kuchanskyy Vladislav, National Academy of Sciences in Ukraine, Ukraine
- Kulakov Andrea, University "Ss. Cyril and Methodius", North Macedonia
- Kulczycki Piotr, Systems Research Institute, Polish Academy of Sciences, Poland
- Kurasova Olga, Institute of Mathematics and Informatics, Bulgaria
- Kusy Maciej, Rzeszow University of Technology, Poland
- Kwasnicka Halina, Wroclaw University of Technology, Poland
- Kwater Tadeusz, PWSTE Jaroslaw, Poland
- Kwolek Bogdan, AGH University of Science and Technology, Poland
- Laccetti Giuliano, University of Naples Federico II and INFN, Italy
- Laghouaouta Youness, INPT, Morocco
- Lakhassane Cissé Mamadou, LIMBI, Senegal
- Lameski Petre, University "Ss.Cyril and Methodius", North Macedonia
- Lasek Piotr, University of Rzeszów, Poland
- Laskov Lasko, New Bulgarian University, Bulgaria
- Lastovetsky Alexey, University College Dublin, Ireland
- Leible Stephan, University of Hamburg, Germany
- Leitao Paulo, Polytechnic Institute of Bragança, Portugal
- Lerga Jonatan, University of Rijeka, Croatia
- Lewandowski Tom, Wrocław University of Science and Technology, Poland
- Lięża Antoni, AGH University of Science and Technology, Poland
- Lilik Ferenc, Szechenyi Istvan University, Hungary
- Lin Zhe, Xiamen University, China
- Ljubić Sandi, University of Rijeka, Faculty of Engineering, Croatia
- Lmati Imane, Faculté Ben Msik, Morocco
- Lobato Fábio, UFOPA, Brasil
- Lovassy Rita, Obuda University, Hungary
- Luque Gabriel, University of Málaga, Spain
- Luszczek Piotr, University of Tennessee, USA
- Machado José, University of Minho, Portugal

- Mačoš Dragan, Beuth Hochschule für Technik, Germany
- Majdik András, HUN-REN SZTAKI - Hungarian Research Network, Hungary
- Malecki Piotr, Institute of Nuclear Physics PAN, Poland
- Mangioni Giuseppe, University of Catania, Italy
- Manso Marco, PARTICLE LTD., Portugal
- Mansurova Madina, al-Farabi Kazakh National University, Kazakhstan
- Marciniak Jacek, Adam Mickiewicz University, Poland
- Marcinkowski Bartosz, University of Gdansk, Poland
- Marcińczuk Michał, Samurai Labs, Poland
- Marghitu Daniela, Auburn University, USA
- Marko Oskar, BioSense Institute, Serbia
- Marowka Ami, Parallel Research Labs, Israel
- Martínez López Pablo E., Universidad Nacional de Quilmes, Argentina
- Masud Mohammad, College of Information Technology, United Arab Emirates
- Matson Eric, Purdue University, USA
- Mazzara Manuel, Innopolis University, Serbia
- Meister Matthias, Intercultural Theological Academy, Liebenzell Mission, Germany
- Mele Valeria, University of Naples Federico II, Italy
- Melzer Sylvia, University of Hamburg, Germany
- Meneses Claudio, Universidad Católica del Norte, Chile
- Mercier-Laurent Eunika, Jean Moulin Lyon 3 University, France
- Merkel Manuel, Universität Stuttgart, Germany
- Mesiar Radko, Slovak University of Technology, Slovakia
- Messe Nan, IRIT, France
- Mhada Fatima Zahra, ENSIAS, Morocco
- Michalik Krzysztof, University of Economics, Katowice, Poland
- Micota Flavia, West University of Timisoara, Romania
- Mignone Paolo, Università degli studi di Bari, Italy
- Mihaescu Marian Cristian, University of Craiova, Romania
- Mihajlov Martin, Jozef Stefan Institute, Slovenia
- Mihálydeák Tamás, University of Debrecen, Hungary
- Milašinović Boris, University of Zagreb, Croatia
- Mildorf Tomas, University of West Bohemia, Czech Republic
- Milella Annalisa, CNR-STIIMA, Italy
- Millham Richard, Durban University of Technology, South Africa
- Milosavljevic Gordana, Faculty of Technical Sciences, Serbia
- Ming Liu David, Khalifa University, United Arab Emirates
- Misra Sanjay, Institute For Energy Technology, Norway
- Mocanu Mihai, University of Craiova, Romania
- Modoni Gianfranco, STIIMA-CNR, Italy
- Mora Andre Damas, UNINOVA, Portugal
- Moroz Leonid, Warsaw University of Technology, Poland
- Moshkov Mikhail, KAUST, Saudi Arabia
- Motii Anas, UM6P, Morocco
- Mozgovoy Maxim, University of Aizu, Japan
- Mullins Roisin, University of Wales Trinity Saint David, United Kingdom
- Mumm Rebekka, WIG2 Institute, Germany
- Munoz Andres, Universidad de Cádiz, Spain
- Muszyńska Karolina, University of Szczecin, Poland
- Myszkowski Pawel, Wrocław University of Science and Technology, Poland
- Nakayama Minoru, Tokyo Institute of Technology, Japan
- Narwal Bhawna, Indira Gandhi Delhi Technical University For Women, India
- Nguyen Hung Son, University of Warsaw, Poland
- Niekler Andreas, Universität Leipzig, Germany
- Niewiadomska-Szynkiewicz Ewa, Warsaw University of Technology, Poland
- Ogrodniczuk Maciej, Polish Academy of Sciences, Poland
- Okarma Krzysztof, West Pomeranian University of Technology in Szczecin, Poland
- Oliveira Eugénio, Universidade do Porto, Portugal
- Oppermann Alexander, Physikalisch-Technische Bundesanstalt, Germany
- Ota Daniel, Fraunhofer, Germany
- Ouariach Soufiane, Abdelmalek Essaadi University, Morocco
- Ozkan Necmettin, Gebze Technical University, Turkey
- Ozkaya Mert, Yeditepe University, Turkey
- Palau Carlos, Universitat Politècnica de Valencia, Spain
- Paliwoda-Pękosz Grażyna, Krakow University of Economics, Poland
- Palma Raul, Poznan Supercomputing and Networking Center, Poland
- Palmigiano Alessandra, the Vrije Universiteit Amsterdam, the Netherlands
- Paluszyński Wiesław, TIC sp. z o.o., Poland
- Pamin Jerzy, Cracow University of Technology, Poland
- Pancerz Krzysztof, The John Paul II Catholic University of Lublin, Poland
- Pandey Dr. Rajiv, Amity University, India
- Pankowska Małgorzata, University of Economics in Katowice, Poland
- Paragliola Giovanni, ICAR-CNR, Italy
- Pataricza András, Budapest University of Technology and Economics, Hungary

- Paziienza Andrea, NTT DATA Italia SpA & A3K Srl, Italy
- Peralta Daniel, Ghent University, Belgium
- Perechuda Kazimierz, Wroclaw University of Economics and Business, Poland
- Peres Ricardo, Instituto de Desenvolvimento de Novas Tecnologia, Portugal
- Petcu Dana, West University of Timisoara, Romania
- Peters Georg, Munich University of Applied Sciences & Australian Catholic University, Germany
- Petrik Milan, Institute of Computer Science (ICS), Czech Republic
- Petrovic Veljko, Faculty of Technical Sciences, Serbia
- Peukert Hagen, University of Hamburg, Germany
- Pinta Pauline Sophia, WIG2 Institute, Germany
- Pirani Massimiliano, Università Politecnica Delle Marche, Italy
- Pires Ivan Miguel, Universidade da Beira Interior, Portugal
- Po Laura, Università di Modena e Reggio Emilia, Italy
- Poczekajło Paweł, Koszalin University of Technology, Poland
- Poledica Ana, University of Belgrade, Serbia
- Porta Marco, University of Pavia, Italy
- Porubán Jaroslav, Technical University of Košice, Slovakia
- Provotar Oleksandr, Taras Shevchenko National University of Kyiv, Ukraine
- Przybyła Piotr, Uniwersytat Pompeu Fabra, Polish Academy of Sciences, Poland
- Przybyła-Kaspepek Małgorzata, Uniwersytet Śląski w Katowicach, Poland
- Ptaszynski Michal, Kitami Institute of Technology, Japan
- Puime Felix, Universidade de A Coruña, Spain
- Rafay Muhammad Abdul, Hasan Murad School of Management (HSM), Pakistan
- Ramanna Sheela, University of Winnipeg, Canada
- Rauch Jan, Prague University of Economics and Business, Czech Republic
- Rechavi Amit, Ruppin Academic Center, Israel
- Reformat Marek, University of Alberta, Canada
- Reis Luis Paulo, Universidade do Porto, Portugal
- Ristic Sonja, University of Novi Sad, Serbia
- Rizvi Syed Tahir Hussain, University of Stavanger, Italy
- Rocha Ana Paula, LIACC, University of Porto, Portugal
- Rocha André, UNINOVA, Portugal
- Rollo Federica, University of Modena and Reggio Emilia, Italy
- Roose Philippe, LIUPPA, France
- Rossi Bruno, Masaryk University, Czech Republic
- Roszczyk Radosław, Warsaw University of Technology, Poland
- Rot Artur, Wroclaw University of Economics, Poland
- Rozevskis Uldis, University of Latvia, Latvia
- Rusho Yonit, Shenkar College of Engineering and Design, Israel
- Ruta Dymitr, Khalifa University, United Arab Emirates
- Sachenko Anatoly, Ternopil State Economic University, Ukraine
- Salem Abdel-Badeeh, Ain Shams University, Egypt
- Salvetti Ovidio, Institute of Information Science and Technologies, National Research Council, Italy
- Samotyj Volodymyr, Lviv State University of Life Safety, Ukraine
- Saraiva Joao, University of Minho, Portugal
- Sarwas Grzegorz, Warsaw University, Poland
- Sawerwain Marek, University of Zielona Góra, Poland
- Schaefer Gerald, Loughborough University, United Kingdom
- Schnepf Timo, Federal Institute for Vocational Education and Training, Germany
- Schreiber Celine, Leipzig University, Germany
- Schreiner Wolfgang, Johannes Kepler University Linz, Austria
- Schreurs Jeanne, Hasselt University, Belgium
- Scozzari Andrea, CNR ISTI, Italy
- Seco Luis, Universidade da Maia, Portugal
- Segedinac Milan, Faculty of Technical Sciences, Novi Sad, Serbia
- Selic Bran, Malina Software Corp., Canada
- Sen Jayanta, Taki Government College, West Bengal, India
- Shah Syed Ihtesham Hussain, ICAR-CNR, Italy
- Sharan Bhagwati, SRM University-AP, Amaravati, India
- Shaska Tony, Oakland University, USA
- Sidje Roger B., University of Alabama, USA
- Siedlecka-Lamch Olga, Czestochowa University of Technology, Poland
- Sierra Jose Luis, Universidad Complutense de Madrid, Spain
- Sifaleras Angelo, University of Macedonia, North Macedonia
- Sikorski Marcin, Gdansk University of Technology, Poland
- Silva Lincoln, UERJ, Brazil
- Simic Dejan, Universität Hamburg, Germany
- Siminski Krzysztof, Silesian University of Technology, Poland
- Singh Pradeep, KIET Group of Institutions, India
- Singh Yashwant, Jaypee University of Information Technology, Wazirpur, India
- Skórzewski Paweł, Adam Mickiewicz University, Poland
- Skubalska-Rafajłowicz Ewa, Wrocław University of Science and Technology, Poland

- Slivnik Boštjan, University of Ljubljana, Slovenia
- Smywiński-Pohl Aleksander, AGH University of Science and Technology, Poland
- Sobczak Andrzej, Warsaw School of Economics, Poland
- Sobińska Małgorzata, Wrocław University of Economics and Business, Poland
- Solanki Vijender Kumar, CMR Institute of Technology, India
- Sorell Thomas, University of Warwick, United Kingdom
- Sousa Pinto Agostinho, CEOS.PP / ISCAP / P-PORTO, Portugal
- Sozer Hasan, Ozyegin University, Turkey
- Stanczyk Urszula, Silesian University of Technology, Poland
- Stankosky Michael, The University of Scranton, USA
- Stark Sandra, Leipzig University, Germany
- Stasolla Fabrizio, Università Giustino Fortunato, Italy
- Stavness Ian, University of Saskatchewan, Canada
- Steinbrink Nicholas, Bertelsmann Stiftung, Germany
- Steiner Petra, BIBB, Germany
- Štěpánek Lubomír, Charles University and Prague University of Economics and Business, Czech Republic
- Stoean Catalin, University of Craiova, Romania
- Stoica Cosmin, University of Craiova, Romania
- Stój Jacek, Silesian University of Technology, Poland
- Stutzer Franziska, WIG2 Institute, Germany
- Subbotin Sergey, National University "Zaporizhzhia Polytechnic", Ukraine
- Suraj Zbigniew, University of Rzeszów, Rzeszów, Poland
- Swacha Jakub, University of Szczecin, Poland
- Symeonidis Symeon, Democritus University of Thrace, Greece
- Szafran Bartłomiej, AGH University of Science and Technology, Poland
- Szantoi Zoltan, European Space Agency, France
- Szczech Izabela, Poznań University of Technology, Poland
- Szczerbicki Edward, University of New Castle, Australia
- Szmit Maciej, University of Łódź, Poland
- Szumski Oskar, University of Warsaw Faculty of Management, Poland
- Szymoniak Sabina, Częstochowa University of Technology, Poland
- Świechowski Maciej, QED Software, Poland
- Taglino Francesco, IASI-CNR, Italy
- Tanwar Sudeep, Nirma University, India
- Tarman Milan, ICS, Slovenija
- Telek Miklos, Budapest University of Technology and Economics, Hungary
- Terra Marcus Vinicius Alencar, Universidade Estadual de Londrina, Brazil
- Tiemann Michael, Federal Institute for Vocational Education and Training, Germany
- Tipparaju Vinod, ByteDance, China
- Tomasz Andrysiak, University of Technology and Life Sciences, Poland
- Tomczyk Łukasz, Jagiellonian University, Poland
- Tomovic Savo, University of Montenegro, Montenegro
- Töreyn Behçet Ugur, Istanbul Technical University, Turkey
- Tormasi Alex, Szechenyi Istvan University, Hungary
- Trentesaux Damien, The Polytechnic University of Hauts-de-France, France
- Trocan Maria, Institut Supérieur d'Électronique de Paris, France
- Tudoroiu Nicolae, John Abbott College, Canada
- Tyagi Sudhanshu, Thapar Institute of Engineering & Technology, India
- Úbeda Ignacio Lacalle, Universitat Politècnica de Valencia, Spain
- Ullah Zaib, Università Telematica Giustino Fortunato, Italy
- Vasilakis Christos, Squaredev, Belgium
- Vasiliev Julian, University of Economics – Varna, Bulgaria
- Vega-Rodríguez Miguel A., University of Extremadura, Spain
- Verstraete Jörg, Instytut Badań Systemowych Polskiej Akademii Nauk, Poland
- Vescoukis Vassilios, National Technical University of Athens, Greece
- Viedma Enrique Herrera, University of Granada, Spain
- Vladusic Daniel, X-LAB, Slovenia
- Vo Bich Khue, University of Finance - Marketing, Vietnam
- Vogiatzis Chrysafis, University of Illinois, USA
- Vu Quang Hieu, Greenfeed, Vietnam
- Wahid Khan Ferdous, Airbus Group, Germany
- Waqas Abdullah, National University of Technology, Pakistan
- Weber Richard, Universidad de Chile, Chile
- Wei Wei, Xi'an University of Technology, China
- Węcel Krzysztof, Poznań University of Economics and Business, Poland
- Wielki Janusz, Opole University of Technology, Poland
- Wimmer Manuel, Johannes Kepler University Linz
- Winnige Stefan, BIBB, Germany
- Wrona Konrad, NATO Communications and Information Agency, the Netherlands
- Wróblewska Anna, Warsaw University of Technology, Poland
- Wyrzykowski Roman, Częstochowa University of Technology, Poland

- Wysocki Marian, Rzeszów University of Technology, Poland
- Xenakis Christos, University of Piraeus, Greece
- Xuetao Jin, Communication University of China, China
- Yang Yujiu, Tsinghua University, China
- Yasir Muhammad, IREA-CNR, Italy
- Zadrożny Sławomir, Systems Research Institute Polish Academy of Sciences, Poland
- Zahra Ouariach Fatima, Abdelmalek Essaadi University Tetouan, Morocco
- Zając Mieczysław, Cracow University of Technology, Poland
- Zborowski Marek, University of Warsaw, Poland
- Zhu Yungang, Jilin University, China
- Zielosko Beata, University of Sielsia, Poland
- Zitouni M. Sami, University of Dubai, United Arab Emirates
- Zukanović Milica, University of Belgrade, Serbia

ACKNOWLEDGMENTS

In conclusion, let us emphasize that delivery of FedCSIS 2024 required a dedicated effort of many people. We would like to express our warmest gratitude to all Topical Area Curators, Thematic Session organizers, members of the FedCSIS 2024 Senior Program Committee and members of the FedCSIS 2024 Program Committee, for their hard work in attracting and reviewing all submissions. We thank the authors of papers for their great contribution to the theory and practice of Computer Science and Intelligence Systems. We are grateful to Keynote and Invited Speakers for sharing their knowledge and experiences with the participants. Last, but not least, we acknowledge, one more time, Ivan Lukovic and Dragana Makajić-Nikolić and their Team, consisting of truly fantastic people. We are very grateful for your efforts!

We hope that you all had an inspiring conference. We also hope to meet you again for the 20th Conference on Computer Science and Intelligence Systems (FedCSIS 2025) which will take place in Kraków, Poland, on September 14-17, 2025. Finally, we hope that you will find the evolution of the FedCSIS Conference concept as something that properly addresses the current needs of research and applications. We want to continue looking at Computer Science from different angles but, at the same time, acknowledging the topic Intelligence Systems as the central point of everything that we are considering (and that has to be considered).

Co-Chairs of the FedCSIS Conference Series:

Marek Bolanowski, Rzeszów University of Technology, Poland

Maria Ganzha, Warsaw University of Technology and Systems Research Institute Polish Academy of Sciences, Poland

Leszek Maciaszek (Honorary Chair), Macquarie University, Australia and Wrocław University of Economics, Poland

Marcin Paprzycki, Systems Research Institute Polish Academy of Sciences and Warsaw University of Management, Poland

Dominik Ślęzak, QED Software and University of Warsaw, Poland

Proceedings of the 19th Conference on Computer Science and Intelligence Systems

September 8–11, 2024. Belgrade, Serbia

TABLE OF CONTENTS

MAIN TRACK

MAIN TRACK INVITED CONTRIBUTIONS

Graded Logic and Professional Decision Making	1
<i>Jozo Dujmović</i>	
Multi-Level Language Architectures as a Foundation for Advanced Enterprise Systems	9
<i>Ulrich Frank</i>	
Machine Learning in Energy and Thermal-aware Resource Management of Cloud Data Centers: A Taxonomy and Future Directions	21
<i>Shashikant Ilager, Rajkumar Buyya</i>	
The Interplay of Learning Analytics and Artificial Intelligence	35
<i>Jelena Jovanovic</i>	
Attentiveness on criticisms and definition about Explainable Artificial Intelligence	45
<i>Francisco Herrera</i>	
How CIs can Tackle Future Pandemics. A Multi-Domain Approach to Improve CI Resilience	53
<i>Stefan Schauer, Manuel Egger, Max Kesselbacher-Pirker, Isti Rodiah, Olga Horvadovska, Berit Lange, Norman FRM Fauster, Hannes Zenz, Christian Kimmich</i>	
Languages for Non-developers: What, How, Where? Invited Talk—Extended Abstract	61
<i>Juha-Pekka Tolvanen</i>	

MAIN TRACK REGULAR PAPERS

Mixed-Methods Study of Arabic Online Review Influence on Purchase Intention (AOCR-PI)	63
<i>Ahmad Alghamdi, Natalia Beloff, Martin White</i>	
Empirical Insights into Cloud Adoption: A new Model Exploring Influencing Factors for Saudi Arabian Small and Medium Enterprises	75
<i>Mohammed Alqahtani, Natalia Beloff, Martin White</i>	
A Quantitative Study Using the ACC-PH Framework: Factors Affecting Cloud Computing Adoption in Saudi Private Hospitals	87
<i>Fayez Alshahrani, Natalia Beloff, Martin White</i>	
A Hybrid Machine Learning Model for Forest Wildfire Detection using Sounds	99
<i>Robertas Damaševičius, Rytis Maskeliunas, Ahmad Qurthobi</i>	
Benchmarking OpenAI's APIs and Large Language Models for Repeatable, Efficient Question Answering Across Multiple Documents	107
<i>Elena Filipovska, Ana Mladenovska, Merxhan Bajrami, Jovana Dobrev, Vellislava Hillman, Petre Lameski, Eftim Zdravevski</i>	

Digital Twin Design for Autonomous Drones	119
<i>Danish Iqbal, Barbora Buhnova</i>	
Critical Success Factors for ERP Projects Revisited: An Update of Literature Reviews	131
<i>Christian Leyh, Alisa Lorenz, Michael Jan Faruga, Linda Koller</i>	
d’Alembert Convolution for Enhanced Spatio-Temporal Analysis of Forest Ecosystems	141
<i>Rytis Maskeliūnas, Robertas Damaševičius</i>	
Optimization of the Cell-based Software Architecture by Applying the Community Detection Approach	149
<i>Miloš Milić, Dragana Makajić-Nikolić</i>	
Model-Agnostic Machine Learning Model Updating – A Case Study on a real-world Application	157
<i>Julia Poray, Bogdan Franczyk, Thomas Heller</i>	
A Blockchain-based Transaction Verification Infrastructure in Public Transportation	169
<i>Hidayet Burak Saritas, Geylani Kardas</i>	
On Privacy of PRF+PUF-based Authentication	177
<i>Ferucio Laurentiu Tiplea</i>	
A Machine Learning Approach for Anxiety and Depression Prediction Using PROMIS Questionnaires	187
<i>Arthur Ricardo Sousa Vitória, Murilo O. Guimarães, Daniel Fazzioni, Aldo A. Díaz-Salazar, Ana Laura S. A. Zara, Iwens G. Sene Junior, Renato F. Bulcão-Neto</i>	
Trust Management Framework for Multi-Robot Systems	195
<i>Daniel Vojnar, Adela Bierska, Barbora Buhnova</i>	
MAIN TRACK SHORT PAPERS	
Enhancing Airbnb Price Predictions with Location-Based Data: A Case Study of Istanbul	207
<i>Özgün Akalın, Gülfem Isiklar Alptekin</i>	
Assessing E-Learning Satisfaction in Saudi Higher Education Post-COVID-19: A Conceptual Framework for e-Services Impact Analysis	213
<i>Wafa Alshammari, Natalia Beloff, Martin White</i>	
LSTM-based Deep Neural Network With A Focus on Sentence Representation for Sequential Sentence Classification in Medical Scientific Abstracts	219
<i>Phat Lam, Lam Pham, Tin Nguyen, Hieu Tang, Michael Seidl, Medina Andresel, Alexander Schindler</i>	
Literature Books Recommender System using Collaborative Filtering and Multi-Source Reviews	225
<i>Elena-Ruxandra Luțan, Costin Bădică</i>	
Toward a Framework for Determining Methods of Evaluation in Design Science Research	231
<i>Julia Müller, Stefanie Würth, Thomas Schäffer, Christian Leyh</i>	
Pathomorphological Diagnosis Process Modeling for Machine Learning Algorithms’ Applying	237
<i>Małgorzata Pańkowska, Mariusz Żytniewski, Mateusz Kozak, Krzysztof Tomaszek, Dominik Spinczyk</i>	
Extension-principle-based Solution Algorithm to Full LR-fuzzy Linear Programming Problems	243
<i>Bogdana Stanojević, Milan Stanojević, Nebojša Nikolić</i>	
Strategy Registry: an optimized implementation of the Strategy design pattern in solidity for the Ethereum Blockchain	249
<i>Hamza Tamenaoul, Mahmoud El Hamlaoui, Mahmoud Nassar</i>	
Successfully Improving the User Experience of an Artificial Intelligence System	253
<i>Alexander Zender, Bernhard G. Humm, Anna Holzheuser</i>	
Task-driven single-image super-resolution reconstruction of document scan	259
<i>Maciej Zyrek, Michal Kawulok</i>	
Congestion Control in Streaming Services with an On-Off MPTCP Algorithm	265
<i>Łukasz Piotr Łuczak, Przemysław Ignaciuk</i>	

THEMATIC SESSIONS

Preface to Thematic Sessions	273
THEMATIC SESSIONS REGULAR PAPERS	
Virtual Power Plant Optimization Service - Benchmark of Solvers	279
<i>Filipe Alves, Rui Ribeiro, Maria Petiz, Ali Abbasi, Pedro Carvalho, Ricardo Faia, Pedro Faria, Zita Vale, Ricardo Rodrigues</i>	
Quality Control of Body Measurement Data Using Linear Regression Methods	289
<i>Janis Bicevskis, Edgars Diebelis, Zane Bicevska, Liva Purina</i>	
Using a Textual DSL With Live Graphical Feedback to Improve the CPS' Design Workflow of Hardware Engineers	301
<i>Twan Bolwerk, Marco Alonso, Mathijs Schuts</i>	
An AI-empowered energy-efficient portable NIRS solution for precision agriculture: A pilot study on a citrus fruit	313
<i>Giulia Cisotto, Dagmawi Delelegn Tegegn, Alberto Zancanaro, Ivan Reguzzoni, Edoardo Lotti, Sara L. Manzoni, Italo F. Zoppis</i>	
Automotive Cybersecurity Engineering with Modeling Support	319
<i>Alexander Fischer, Juha-Pekka Tolvanen, Ramin Tavakoli Kolagari</i>	
Goliath, a Programming Exercises Generator Supported by AI	331
<i>Tiago Carvalho Freitas, Alvaro Costa Neto, Maria João Varanda Pereira, Pedro Rangel Henriques</i>	
IoB-TMAF: Internet of Body-based Telemedicine Adoption Framework	343
<i>Taif Ghiwaa, Imran Khan, Martin White, Natalia Beloff</i>	
Linked Labor Market Data: Towards a novel data housing strategy	355
<i>Kristine Hein</i>	
An autonomous vehicle in a connected environment: case study of cyber-resilience	363
<i>Guillaume Hutzler, Hanna Klaudel, Witold Klaudel, Franck Pommereau, Artur Rataj</i>	
Learning from the COVID-19 Pandemic to Improve Critical Infrastructure Resilience using Temporal Fusion Transformers	375
<i>Jakob Jenko, Joao Pita Costa, Daniel Vladušič, Urban Bavčar, Radoš Čabarkapa</i>	
HPC operation with time-dependent cluster-wide power capping	385
<i>Alexander Kammeyer, Florian Burger, Daniel Lübbert, Katinka Wolter</i>	
Teaching Beginners to Program: should we start with block-based, text-based, or both notations?	395
<i>Tomaž Kosar, Srđa Bjeladinović, Dragana Ostojić, Milica Škembarević, Žiga Leber, Olga Jejić, Filip Furtula, Miloš Ljubisavljević, Ivan Luković, Marjan Mernik</i>	
Pareto Optimal Solutions of the Biobjective Minimum Length Minimum Risk Spanning Trees Problem	405
<i>Lasko Laskov, Marin Marinov</i>	
Towards the analysis of errors in centrality measures in perpetuated networks	417
<i>Meetkumar Pravinbhai Mangroliya, Jens Doerpinghaus, Robert Rockenfeller</i>	
Real options analysis framework for agile projects	429
<i>Gloria J. Miller</i>	
MBSPI - A Model-Based Security Pattern Integration Approach for software architectures	443
<i>Anas Motii, Mahmoud El Hamlaoui</i>	
An Ontology to Understand Programming Cocktails	453
<i>Alvaro Costa Neto, Maria João Varanda Pereira, Pedro Rangel Henriques</i>	
SrpCNNeL: Serbian Model for Named Entity Linking	465
<i>Milica Ikonić Nešić, Saša Petalinkar, Ranka Stanković, Miloš Utvić, Olivera Kitanović</i>	
Towards crop traits estimation from hyperspectral data: evaluation of neural network models trained with real multi-site data or synthetic RTM simulations	475
<i>Lorenzo Parigi, Gabriele Candiani, Ignazio Gallo, Piero Toscano, Mirco Boschetti</i>	

AI-Based Spatiotemporal Crop Monitoring by Cloud Removal in Satellite Images	485
<i>Jiří Pihrt, Petr Šimánek, Alexander Kovalenko, Jiří Kvapil, Karel Charvát</i>	
Organizational Capabilities for Business-IT Integration in Digital Enterprises	493
<i>Constanze Riedinger, Maike Netscher, Stephan Zimmermann</i>	
Automated feedback generation in an intelligent tutoring system for counselor education	501
<i>Eric Rudolph, Hanna Seer, Carina Mothes, Jens Albrecht</i>	
A framework for enabling ex-ante social impact assessment of project-based technological solutions: the case of Remote Infrastructure Inspection	513
<i>Nikolay Zherdev, Olivier Klein, Umberto Sconfienza, Philippe Gerber, Daniel Vladušič, Jethro Butler, Aljosa Pasic</i>	
THEMATIC SESSIONS SHORT PAPERS	
LeAF: Leveraging Deep Learning for Agricultural Pest Detection and Classification for Farmers	525
<i>Aditya Sengupta</i>	
Key Factors Influencing Mobile Banking Adoption in Saudi Arabia	531
<i>Amal Alzahrani, Natalia Beloff, Martin White</i>	
A novel ensemble learning technique of shallow models applied on a COVID-19 dataset	537
<i>Diogen Babuc</i>	
Agricultural Data Space: the METRIQA Platform and a Case Study in the CODECS project	543
<i>Manlio Bacco, Alexander Kocian, Antonino Crivello, Marco Gori, Giovanna Maria Dimitri, Paolo Barsocchi, Gianluca Brunori, Stefano Chessa</i>	
Reinforcement Learning based Intelligent System for Personalized Exam Schedule	549
<i>Marco Barone, Matteo Ciaschi, Zaib Ullah, Armando Piccardi</i>	
Disease Diagnosis On Ships Using Hierarchical Reinforcement Learning	555
<i>Farwa Batool, Tehreem Hasan, Giancarlo Tretola, Zaib Ullah, Musarat Abbas</i>	
Dashboard User interface (UI) Implementation for remote critical infrastructure inspection by using UAV/Satellite in times of pandemic	561
<i>Romaio Bratskas, Dimitrios Papachristos, Petros Savvidis, George Leventakis, Enea Qerama, George Dahrouje</i>	
Evolving the Enterprise Software Systems Landscape: Towards Anti-Patterns in Smalltalk-to-Java Code Transformation	567
<i>Marek Bělohoubek, Robert Pergl</i>	
Exploring the role of Artificial Intelligence in assessing soft skills	573
<i>Matteo Ciaschi, Marco Barone</i>	
Spoken Language Corpora Augmentation with Domain-Specific Voice-Cloned Speech	579
<i>Mateusz Czyżnikiewicz, Łukasz Bondaruk, Jakub Kubiak, Adam Wiącek, Łukasz Degórski, Marek Kubis, Paweł Skórzewski</i>	
Comparing Lazy Constraint Selection Strategies in Train Routing with Moving Block Control	585
<i>Stefan Engels, Robert Wille</i>	
Dynamic Threat Intelligence for Improvement of Resilience of Critical Infrastructure During Pandemics	591
<i>Pablo de Juan Fidalgo, Aljosa Pasic, Susana González Zarzosa</i>	
Hospital Patient Distribution After Earthquake	597
<i>Stefka Fidanova, Leoneed Kirilov, Veselin Ivanov, Maria Ganzha</i>	
Impact of Spelling and Editing Correctness on Detection of LLM-Generated Emails	603
<i>Paweł Gryka, Kacper Gradoń, Marek Kozłowski, Miłosz Kutyla, Artur Janicki</i>	
A network clustering method based on intersection of random spanning trees	609
<i>László Hajdu, András London, András Pluhár</i>	
Efficient Maritime Healthcare Resource Allocation Using Reinforcement Learning	615
<i>Tehreem Hasan, Farwa Batool, Mario Fiorino, Giancarlo Tretola, Musarat Abbas</i>	

Unconditional Token Forcing: Extracting Text Hidden Within LLM	621
<i>Jakub Hořcitolowicz, Paweł Popiołek, Jan Rudkowski, Jędrzej Bieniasz, Artur Janicki</i>	
Plant-traits: how citizen science and artificial intelligence can impact natural science	625
<i>Giacomo Ignesti, Davie Moroni, Massimo Martinelli</i>	
Agent at the Edge: Opportunity and Challenges of Video Streaming Analytics at the CDN Edge	631
<i>Reza Shokri Kalan, Seren Gul</i>	
DSML4JaCaMo: A Modelling tool for Multi-agent Programming with JaCaMo	637
<i>Burak Karaduman, Baris Tekin Tezel, Geylani Kardas, Moharram Challenger</i>	
Towards understanding animal welfare by observing collective flock behaviors via AI-powered Analytics	643
<i>Savvas Karatsiolis, Pieris Panagi, Vassilis Vassiliades, Andreas Kamilaris, Nicolas Nicolaou, Efsthios Stavrakis</i>	
Topic Modeling of the SrpELTeC Corpus: A Comparison of NMF, LDA, and BERTopic	649
<i>Teodora Mihajlov, Milica Ikonić Nešić, Ranka Stanković, Olivera Kitanović</i>	
Towards to an interface design for a building operations CPS	655
<i>Filipe Moreira, Rosana Alexandre, Rosa Mariana Silva, João Oliveira, Manuel Alves, João Pereira, Ana Colim, Nelson Rodrigues</i>	
An environment model in multi-agent reinforcement learning with decentralized training	661
<i>Rafał Niedziółka-Domański, Jarosław Bylina</i>	
Lower Bounds on Cardinality of Reducts for Decision Tables from Closed Classes	667
<i>Azimkhon Ostonov, Mikhail Moshkov</i>	
Automatic Generation of OpenCL Code through Polyhedral Compilation with LLM	671
<i>Marek Palkowski, Mateusz Grużewski</i>	
Impact of Local Geometry on Methods for Constructing Protein Conformations	677
<i>Wagner Da Rocha, Therese Malliavin, Antonio Mucherino, Leo Liberti</i>	
The comparison of pixel-based image analysis for detection of weeds in winter wheat from UAV imagery	683
<i>Vojtěch Slezák, Kateřina Kuchaříková, Tomáš Kaplánek, Vojtěch Lukas, Jan Křen</i>	
An Improved Genetic Algorithm for Set Cover using Rosenthal Potential	689
<i>Dena Tayebi, Saurabh Ray, Deepak Ajwani</i>	
Efficiency and Reliability of Avalanche Consensus Protocol in Vehicular Communication Networks	695
<i>Saeed Ullah, Zaib Ullah, Abdullah Waqas</i>	
Stacking Ensemble Machine Learning Modelling for Milk Yield Prediction Based on Biological Characteristics and Feeding Strategies	701
<i>Ruiming Xing, Baihua Li, Shirin Dora, Michael Whittaker, Janette Mathie</i>	
A bottom-up approach to select constrained spectral bands discriminating vine diseases	707
<i>Shurong Zhang, Alban Goupil, Valeriu Vrabie, Eric Perrin, Marie-Laure Panon</i>	
Dynamic relationship between population densities and air quality in the four largest Norwegian cities	713
<i>Petar Zhivkov, Todor Kesarovski</i>	
Secretary problem revisited: Optimal selection strategy for top candidates using one try in a generalized version of the problem	719
<i>Lubomír Štěpánek</i>	
Non-parametric comparison of survival functions with censored data: A computational analysis of greedy and Monte Carlo approaches	725
<i>Lubomír Štěpánek, Filip Habarta, Ivana Malá, Luboš Marek</i>	

COMPETITIONS

DATA MINING COMPETITION

Predicting Stock Trends Using Common Financial Indicators: A Summary of FedCSIS 2024 Data Science Challenge Held on KnowledgePit.ai Platform	731
<i>Aleksandar M. Rakićević, Pavle D. Milosević, Ivana T. Dragović, Ana M. Poledica, Milica M. Zukanović, Andrzej Janusz, Dominik Ślęzak</i>	
Decoding Financial Data: Machine Learning Approach to Predict Trading Actions	739
<i>Yat Chun Fung, Bekzod Amonov</i>	
Searching Stable Solutions For Stock Predictions: A Stacking Approach	745
<i>Ty Gross, Arthur Allebrandt Werlang, Apeksha Poudel, Julian Roß</i>	
Experimenting with manual and automated data mining pipelines on the FedCSIS 2024 Data Science Challenge	751
<i>Max Lautenbach, Jusstina Judák, Luisa Buck, Marc Furier, Okan Mert Göktepe, Gregor Münker</i>	
Key Financial Indicators Analysis and Stock Trend Forecasting Based on a Wrapper Feature Selection Method	755
<i>Chang Lin</i>	
Exploring Stability and Performance of hybrid Gradient Boosting Classification and Regression Models in Sectors Stock Trend Prediction: A Tale of Preliminary Success and Final Challenge	761
<i>Ming Liu, Ling Cen, Dymitr Ruta, Quang Hieu Vu</i>	
Forecasting Stock Trends with Feedforward Neural Networks	767
<i>Marcin Traskowski, Eyad Kannout</i>	
Author Index	773

Graded Logic and Professional Decision Making

Jozo Dujmović
0000-0002-1715-2700
San Francisco State University
Department of Computer Science
1600 Holloway Ave., San
Francisco, CA 94132, USA
jozo@sfsu.edu

Abstract—This paper summarizes the basic concepts of Graded Logic (GL) and the use of GL in professional decision making. Our goal is to contrast two approaches to the development of a continuum-valued propositional logic: (1) the human-centric approach based on observing and modeling human commonsense logical reasoning in the context of decision making, and (2) the theoretical approach where logic is developed as a formal axiomatic deductive system. We show the basic advantages of human-centric approach and the applicability of this approach in the area of professional decision making.

Index Terms—Graded Logic, commonsense logical reasoning, decision making, LSP method, GCD function.

I. INTRODUCTION

LOGIC is a wide area studied in both philosophy and mathematics [1]. In this short survey paper, we are interested only in the propositional logic [2], i.e., the logic that combines degrees of truth of input statements to compute the degree of truth of a compound output statement. We briefly present a human-centric Graded Logic (GL) which is derived from observing, measuring and modeling human commonsense logical reasoning in the process of decision making. We also present a brief survey of the Logic Scoring of Preference (LSP) decision method [3] which is based on Graded Logic [22].

The classical logic [4]-[6] and its modern extensions [7], [8], as well as non-classical logics [9], [10] are created as formal axiomatic deductive systems; that is the standard theoretical approach. In such systems a mathematical theory is built on a set of axioms and axioms are assumed to be true without further consideration. Then, all other theoretical results are proved based on their consistency with the axioms. In the case of logic, axioms operate with variables that are members of a set of two or more values, but further considerations of the role and meaning of variables are not necessary and not given. So, such logics operate with anonymous real numbers, and if such numbers denote degrees of truth of specific statements, it is not necessary to know the corresponding statements, their

author, and their role, meaning, and the context in which the statements are created and used. The applicability of axiomatic deductive logic systems is an independent topic outside the area of specific (logic) theory.

It is possible to develop logic in a different, human-centric way, which we propose in the case of GL. This approach is derived from logic-based applications of specific stakeholder/decision-maker who is an individual human or a human organization. The human-centric approach is based on observing, measuring, modeling and explaining natural human commonsense reasoning and decision making. Mathematical models are then developed to be consistent with observations and measurements.

The paper is organized as three sections devoted to Graded Logic, followed by a section devoted to the LSP method.

II. HUMAN-CENTRIC APPROACH TO LOGIC

A. The stakeholder/decision-maker

Logical reasoning is a human mental activity, i.e., there is no logical reasoning and no need for logic without explicit presence of a specific human (either an individual or an organization). Human logic does not exist in vacuum. Therefore, we assume that all logic problems are related to a specific human participant, identified as the stakeholder/decision-maker (SDM).

It should be self-evident that the SDM exists in a specific environment, interacts with the environment, has goals and requirements, and uses logical reasoning to make decisions necessary to satisfy requirements and attain goals.

B. Human graded percepts and graded truth

Human percepts are defined as quantifiable mental sensations/impressions of perceiving and/or reasoning. Examples of such percepts include satisfaction, importance, suitability, preference, confidence, value, and many others. The fundamental property of such percepts is that they are graded: each percept p can vary in the range from zero to its maximum

value: $p \in [0, p_{max}]$. For example, each percept of satisfaction of specific requirements varies in the range from no satisfaction to the full satisfaction.

All graded percepts can be directly related to graded truth. If we define $t = p/p_{max}$, $t \in [0,1]$, then t denotes the degree of truth of statement “the percept p attained its maximum value.” For example, if p is a percept of satisfaction with a family car, then t is the degree of truth of statement “the car fully satisfies all our requirements.” Obviously, if the car satisfies all requirements only in 70% of cases, then $t = 0.7$ and t is a continuum-valued graded truth. Graded Logic is a propositional calculus that processes graded truth.

C. Graded Logic and decision making

Decision making is an observable human mental process based on commonsense logical reasoning. In the most frequent case, the SDM first identifies a set of m different alternatives that can be applied to attain desired goals. The decision making can then be defined as the process of comparison of alternatives and selection (and possible realization) of the best alternative. To understand the process of human decision making and its relation to GL, there is a prerequisite: it is necessary to understand the fundamental case $m = 1$.

D. The case of single alternative

The case of a single alternative is not a special case. In decision making, that is the most important essential case. It is easy to find a number of single alternative decision problems in each human life. E.g., the most important decision in most human lives is marriage, and there is a single candidate that must be carefully evaluated. Indeed, the question is how suitable a single candidate is, and not who is the best (i.e. the least unsuitable) among several candidates. The best among several candidates/options, selected using pairwise comparison methods, can still be insufficiently suitable and justifiably rejected.

Similarly, an unemployed worker can get a single job offer and the question is whether the offer is sufficiently good to be accepted. On the other hand, a company can have a single candidate applying for an open position, and it is necessary to evaluate the competence of the single candidate and then either to accept or to reject the candidate.

The presented examples expose the evaluation process of a single candidate as a fundamental component of human-centric decision making. If a single candidate evaluation process is available, then the comparison of multiple candidates is automatically solved by comparing the results of evaluation of individual candidates.

E. The commonsense evaluation process and its logic components

The natural human commonsense evaluation process has the following easily visible components [22]:

1. Selection of suitability attributes.
2. Development of suitability attribute criteria.
3. Generating the attribute suitability degrees.
4. Logic aggregation of attribute suitability degrees.

5. Evaluation and comparison of alternatives.

We assume that the SDM has clearly defined goals and can specify requirements that the available objects/alternatives should satisfy. The first step performed by SDM is the selection of suitability attributes. Suitability attributes are all those characteristics of the evaluated objects that affect the overall suitability of each evaluated object. For example, if the evaluated object is a car, then the suitability attributes could include the power of engine, fuel economy, available space, the number of passengers, wheel drive, etc. It is also important to note that there are attributes of the evaluated object that according to the SDM's goals do not affect the suitability of evaluated objects/alternatives and such attributes are not considered by the SDM. Suitability attributes are denoted a_1, \dots, a_n , $a_i \in \mathbb{R}$, $i = 1, \dots, n$. Regularly, $n > 1$.

The second step is the definition of requirements that the suitability attributes must satisfy according to SDM's goals and needs. Such requirements are the attribute criteria, i.e. functions that specify the way SDM determines the suitability of each attribute. The suitability is a graded percept expressed as the graded truth of the statement that asserts the complete/full satisfaction of SDM's requirements. So, the attribute criteria are $g_i: \mathbb{R} \rightarrow [0,1]$, $i = 1, \dots, n$.

In the third step, the SDM separately evaluates each attribute of an evaluated object/alternative, and creates the attribute suitability degrees which are percepts of satisfaction of requirements each attribute is expected to satisfy. So, the SDM intuitively creates n degrees of truth $x_i = g_i(a_i)$, $i = 1, \dots, n$.

The availability of n individual percepts of satisfaction of the suitability attributes requirements is the result of the three initial steps of the commonsense evaluation. In the fourth step of human commonsense decision making, the individual percepts x_1, \dots, x_n automatically contribute to forming a resulting graded percept of the overall satisfaction of requirements, $X = L(x_1, \dots, x_n)$. The aggregation function $L: [0,1]^n \rightarrow [0,1]$ is obviously a propositional logic formula. This function combines the models of simultaneity (graded conjunction), substitutability (graded disjunction) and negation. Such combinations of basic graded logic functions generate a wide spectrum of commonsense propositional calculus logic models used in human commonsense reasoning and decision making.

In the fifth (final) step of commonsense decision making, the overall satisfaction of requirements $X \in [0,1]$ is used to decide whether to accept or to reject a specific object/alternative. In addition, the degrees of overall suitability are used in the process of explainable commonsense comparison and selection of multiple competitive objects/alternatives.

III. THE CONCEPT OF FULLY CONTINUUM-VALUED LOGIC

A. The continuum-valued graded percepts

Human commonsense logical reasoning is based on graded percepts. The primary graded percept is the graded truth. It specifies the intensity of a specific graded percept as the degree of truth of a statement that claims the highest level of the

percept. For example, if a car engine should ideally have 200 HP, then the degree of truth of the assertion that the car engine of 180 HP fully satisfies SDM's requirements could be 0.9.

Truth is not the only continuum-valued human graded percept. The second fundamental graded percept is importance. It is easy to note that in human commonsense logical reasoning some statements are more important than other statements. For example, for computationally intensive tasks, the processor speed of a laptop computer can be significantly more important than the weight of computer. So, the importance of statements aggregated by a graded logic function must also be continuum-valued. Human commonsense logical reasoning supports the "first things first" concept.

The most distinctive property of both the commonsense human logic and the Graded Logic is that both simultaneity and substitutability are *graded*: their intensity is continuously adjustable. Below, we discuss this characteristic property.

B. Unification of simultaneity and substitutability

In human commonsense logic, the simultaneity (graded conjunction) and the substitutability (graded disjunction) are not treated as two separated and different logic operations. Each human logic aggregator of two or more variables has both conjunctive and disjunctive properties. Conjunctive properties in evaluation are specified as a requirement that all inputs should simultaneously have (to some desired extent) high values. On the other hand, required disjunctive properties mean that a low satisfaction of any input can (to some desired extent) be substituted/compensated by a high value of any other input. These opposing requirements can be balanced in the case of the arithmetic mean, where the conjunctive properties are equally present as disjunctive properties. A typical example is the computation of the mean grade of students in schools (GPA), where high grades are simultaneously desired in all courses, but at the same time, a low grade in any course can be compensated by a high grade in any other course.

In Graded Logic [3], the logic aggregator that combines conjunctive and disjunctive properties is called Graded Conjunction/Disjunction (GCD) and denoted $y = x_1 \diamond \dots \diamond x_k$ (the symbol \diamond is a combination of symbols \wedge and \vee).

C. Andness and orness

Simultaneity and substitutability are graded, i.e., they also have adjustable intensity. In the case of simultaneity, the SDM may want that two (or more) requirements are simultaneously highly satisfied (e.g. all product buyers simultaneously want a high quality *and* a low price of selected product). It is easy to note that, in human reasoning, the intensity of simultaneity for conjunctive logic aggregators is continuously adjustable. In the case of high degree of simultaneity, SDMs frequently use *mandatory requirements*: if one of inputs is not satisfied, then the results of aggregation must be zero (i.e., such a function supports the annihilator 0). In the case of low intensity, the simultaneous satisfaction of inputs can be desirable but not mandatory. In such cases, the annihilator 0 must not be supported.

In the case of substitutability, the situation is similar: the disjunctive aggregators also have an adjustable intensity. High intensity disjunctive aggregators support the annihilator 1: if any of inputs is fully satisfied, then the high-intensity disjunctive aggregator is fully satisfied. In the case of lower intensity, the high degree of satisfaction of inputs is desirable, but not individually sufficient to fully satisfy a disjunctive criterion.

The intensity of simultaneity is called the conjunction degree or andness [11], [12] and denoted α . The intensity of substitutability is called the disjunction degree or orness and denoted ω . For the GCD aggregator $y = x_1 \diamond \dots \diamond x_k$ they are defined as follows [11]:

$$\alpha = \frac{k}{k-1} - \frac{k+1}{k-1} \int_{[0,1]^k} (x_1 \diamond \dots \diamond x_k) dx_1 \dots dx_k$$

$$\omega = 1 - \alpha .$$

According to [13],

$$\int_{[0,1]^k} (x_1 \wedge \dots \wedge x_k) dx_1 \dots dx_k = \frac{1}{k+1} ,$$

$$\int_{[0,1]^k} (x_1 \vee \dots \vee x_k) dx_1 \dots dx_k = \frac{k}{k+1} .$$

Therefore, for $x_1 \wedge \dots \wedge x_k$ we have $\alpha = 1$, $\omega = 0$, and for $x_1 \vee \dots \vee x_k$ we have $\alpha = 0$, $\omega = 1$. Another important case is $x_1 \diamond \dots \diamond x_k = (x_1 + \dots + x_k)/k$ where $\alpha = \omega = 1/2$. Therefore, the arithmetic mean has the central logically neutral role as the centroid of GCD logic aggregators.

D. Duality of simultaneity and substitutability

Duality of simultaneity and substitutability is a natural property of commonsense human logic. Let $x_1 \Delta \dots \Delta x_k$ denote the simultaneity (graded conjunction) of k logic variables and let $x_1 \nabla \dots \nabla x_k$ denote the substitutability (graded disjunction) of the same logic variables. Then, the commonsense verbal interpretation of relationship between graded conjunction and graded disjunction is "if we need simultaneously high satisfaction (truth values) of k inputs, then it is not acceptable that any one of them is not sufficiently satisfied." In other words, $x_1 \Delta \dots \Delta x_k = 1 - (1 - x_1) \nabla \dots \nabla (1 - x_k)$. Similarly, "if we need at least one sufficiently satisfied input, then it is not acceptable that all of them are simultaneously insufficiently satisfied." Thus, $x_1 \nabla \dots \nabla x_k = 1 - (1 - x_1) \Delta \dots \Delta (1 - x_k)$. Therefore, it is obvious that De Morgan duality is naturally present in the commonsense human logic.

In the duality relationships we assume the same intensity of conjunctive and disjunctive aggregators Δ and ∇ . If they have the highest idempotent intensity $\Delta = \wedge$ and $\nabla = \vee$ then we get the traditional De Morgan laws:

$$x_1 \wedge \dots \wedge x_k = \overline{\overline{x_1} \vee \dots \vee \overline{x_k}} ; x_1 \vee \dots \vee x_k = \overline{\overline{x_1} \wedge \dots \wedge \overline{x_k}} .$$

In the case of low intensity where $\Delta = \nabla = \diamond$ (the symbol \diamond denotes the arithmetic mean $x_1 \diamond \dots \diamond x_k = w_1 x_1 + \dots + w_k x_k$, $0 < w_i < 1$, $i = 1, \dots, k$, $w_1 + \dots + w_k = 1$) we have

$$1 - (1 - x_1) \diamond \dots \diamond (1 - x_k) = 1 - [w_1(1 - x_1) + \dots + w_k(1 - x_k)] = w_1 x_1 + \dots + w_k x_k .$$

In Graded Logic duality holds for all logic aggregators: soft idempotent, hard idempotent, and nonidempotent hard hyperconjunction and hyperdisjunction [3].

E. The drastic conjunction

What is the strongest possible conjunction? The highest level of simultaneity of high values of k inputs is obviously the extreme requirement that all inputs must be fully satisfied $x_1 = \dots = x_k = 1$. In all other cases the result of conjunctive aggregation is 0. Such a conjunctive function is called drastic conjunction and its analytic form is $y = [\prod_{i=1}^k x_i]$. Since $\int_{[0,1]^k} [\prod_{i=1}^k x_i] dx_1 \dots dx_k = 0$, it follows that the highest possible andness is $\alpha = \frac{k}{k-1}$. The lowest possible orness for drastic conjunction is $\omega = 1 - \alpha = \frac{-1}{k-1}$.

F. The drastic disjunction

The drastic disjunction is a function that is the De Morgan dual of drastic conjunction: $y = 1 - [\prod_{i=1}^k (1 - x_i)]$. So, the strongest possible disjunction is the case where any nonzero input can fully satisfy the disjunctive criterion which is not satisfied if and only if all inputs are zero. Then, we have $\int_{[0,1]^k} \{1 - [\prod_{i=1}^k (1 - x_i)]\} dx_1 \dots dx_k = 1$ and therefore, $\alpha = -\frac{1}{k-1}$, and $\omega = \frac{k}{k-1}$.

G. The interpolative GCD logic aggregator

The extreme drastic conjunction and drastic disjunction aggregators show that the GCD aggregator must cover the full range of andness and orness $[-\frac{1}{k-1}, \frac{k}{k-1}]$. To provide a continuous transition in this wide range of andness, we use interpolative logic aggregators [14]. We select a sequence of conjunctive ‘‘anchor aggregators:’’

- Logic neutrality (arithmetic mean, $\alpha = 0.5$)
- Threshold hard conjunction ($\alpha = 0.75$)
- Pure conjunction ($\alpha = 1$)
- Product t-norm ($\alpha = (k2^k - k - 1)/(k - 1)2^k$)
- Drastic conjunction ($\alpha = k/(k - 1)$)

Between the anchor aggregators with andness α_p and α_q we use interpolation:

$$GCD(\mathbf{x}; \alpha) = \frac{\alpha_q - \alpha}{\alpha_q - \alpha_p} GCD(\mathbf{x}; \alpha_p) + \frac{\alpha - \alpha_p}{\alpha_q - \alpha_p} GCD(\mathbf{x}; \alpha_q)$$

$$\alpha_p \leq \alpha \leq \alpha_q, \quad \mathbf{x} = (x_1, \dots, x_k).$$

Between the logic neutrality and threshold conjunction the interpolated GCD aggregators are soft (the annihilator 0 is not supported). Above the threshold andness the interpolated GCD aggregators are hard (the annihilator 0 is supported). Below the pure conjunction the GCD aggregators are idempotent and above the pure conjunction they are nonidempotent. We use this interpolative form of GCD for $\alpha \geq 0.5$. In the disjunctive range of andness ($\alpha < 0.5$) we recursively use De Morgan duals of the conjunctive GCD:

$$GCD(\mathbf{x}; \alpha) = 1 - GCD(\mathbf{1} - \mathbf{x}; 1 - \alpha), \quad \alpha < \frac{1}{2}.$$

H. Fully continuum-valued propositional logic

Graded Logic is a fully continuum-valued propositional logic of human commonsense reasoning. Same as in natural

human commonsense reasoning, everything is a matter of degree: truth, importance, conjunction (simultaneity), and disjunction (substitutability) are continuum-valued (graded).

The concept of making a propositional logic consistent with natural commonsense human reasoning is easily justifiable by the fact that decision making is a human mental activity. Logic models that are not consistent with observable properties of human reasoning cannot generate results that are explainable and acceptable with confidence. Indeed, the credibility of decision methods that are not consistent with human commonsense decision making is generally questionable.

IV. MAIN PROPERTIES OF GRADED LOGIC

A. The postulates of Graded Logic

Graded Logic is not a formalized axiomatic theory, but it is built on a set of strict postulates that reflect the observable properties of human commonsense logic. There are ten such postulates [22]:

1. The truth of statements must be continuum-valued (graded).
2. The importance of statements must be continuum-valued (graded).
3. The simultaneity of statements must be continuum-valued (graded) in the full range $\frac{1}{2} \leq \alpha \leq \frac{k}{k-1}$.
4. The substitutability of statements must be continuum-valued (graded) in the full range $\frac{1}{2} \leq \omega \leq \frac{k}{k-1}$.
5. The simultaneity and substitutability must be complementary and unified.
6. Logic neutrality (arithmetic mean) must be available as a balance of simultaneity and substitutability.
7. The idempotency of logic aggregators must be selectable.
8. The annihilator support for idempotent simultaneity must be selectable.
9. The annihilator support for idempotent substitutability must be selectable.
10. The simultaneity and substitutability models must be dual.

B. The Graded Logic Conjecture

Each propositional calculus uses a set of basic logic functions to create compound logic formulas. GL is a propositional logic and therefore the fundamental question is to select the necessary and sufficient basic logic functions of graded propositional calculus. According to *Graded Logic Conjecture (GLC)* [3], necessary and sufficient basic Graded Logic functions include ten characteristic functions: nine characteristic special cases of the GCD aggregator and negation. Following are the GLC functions, classified by andness and their support of GL postulates [22]:

1. Graded hyperconjunction ($\alpha > 1$) [C/A0/NI]
2. Pure conjunction – minimum ($\alpha = 1$) [C/A0/ID]
3. Hard graded conjunction ($0.75 \leq \alpha < 1$) [C/A0/ID]
4. Soft graded conjunction ($0.5 < \alpha < 0.75$) [C/NA/ID]

- | | |
|---|-----------|
| 5. Logic neutrality ($\alpha = 0.5$) | [N/NA/ID] |
| 6. Soft graded disjunction ($0.25 < \alpha < 0.5$) | [D/NA/ID] |
| 7. Hard graded disjunction ($0 < \alpha \leq 0.25$) | [D/A1/ID] |
| 8. Pure disjunction - maximum ($\alpha = 0$) | [D/A1/ID] |
| 9. Graded hyperdisjunction ($\alpha < 0$) | [D/A1/NI] |
| 10. Negation (which is not an aggregator) | |

The classification codes [type/annihilators/idempotence] are the following: C = conjunctive, D = disjunctive, N = neutral, A0 = supports annihilator 0, A1 = supports annihilator 1, NA = no support for annihilators, ID = idempotent, NI = nonidempotent.

The GLC is supported by the following properties/facts:

- The GLC functions explicitly support the functionality requested in the postulates of Graded Logic
- All ten GLC functions are provably used in human commonsense logical reasoning.
- Observations of human commonsense logical reasoning have not detected logical reasoning patterns that are not modellable by the presented GLC list of basic GL functions. In particular, all canonical logic aggregation structures detected in the area of decision making, and reported in [3], are modellable using combinations of GLC functions.
- The GLC functions include pure conjunction, disjunction and negation which are necessary and sufficient in classical Boolean logic, making GL a generalization of the classical Boolean logic.

Since GL is not an axiomatic formal system, the satisfaction of GL postulates and consistency with observable commonsense human logic can be used as a sufficient support for conclusion that ten GLC functions are necessary and sufficient to create all formulas of the graded propositional calculus used in natural human reasoning. Hyperconjunction, hard and soft conjunctive and disjunctive GCD, and hyperdisjunction have continuously adjustable andness/orness in their respective ranges.

GCD functions that support annihilators (A0, A1) are denoted as *hard*, and GCD functions that do not support annihilators (NA) are denoted as *soft*. These logic aggregators have the following verbalized interpretation:

- *Must have all inputs:* hard conjunctive.
- *Nice to have most inputs:* soft conjunctive.
- *Nice to have some inputs:* soft disjunctive.
- *Enough to have any input:* hard disjunctive.

According to interpolative method of GCD design, the anchor aggregators have a constant andness, and the inter-anchor aggregators cover a range of andness. The locations of anchor aggregators that are thresholds between soft and hard GCD aggregators are freely adjustable, but in [14] the uniform distribution of hard and soft properties based on thresholds $\alpha = 0.75$ and $\alpha = 0.25$ is experimentally verified to be the closest to the commonsense human logical reasoning. The properties of the GCD function in the full range of andness (from the drastic conjunction CC to the drastic disjunction

DD) are shown in Fig. 1 (dark gray area = soft conjunction/disjunction, light gray area = hard conjunction/disjunction, white area = hyperconjunction and hyperdisjunction).




Fig 1. GCD logic aggregator in the full range of andness/orness

C. Advanced Graded Logic constructs

Combinations of GLC functions yield an infinite number of possible propositional calculus formulas. Some of them have a frequent use in logic decision models. Three most important constructs are (1) the partial absorption, (2) the selector-based formulas, and (3) the nonstationary formulas.

The *partial absorption function* [12], [15], [16], [3], [22] aggregates two asymmetric inputs: mandatory/optional and sufficient/optional. It uses the weighted arithmetic mean ϕ , the hard partial conjunction Δ , and the hard partial disjunction ∇ to create the *conjunctive partial absorption function* $CPA(x, y) = x\Delta(x \phi y)$ and the *disjunctive partial absorption function* $DPA(x, y) = x\nabla(x \phi y)$. The basic property of CPA is $CPA(0, y) = 0$, showing that x is a mandatory input that must be satisfied. The basic property of DPA is $DPA(1, y) = 1$, showing that x is a sufficient input and it is enough to fully satisfy this input. In both cases, y is an optional input. If $0 < x < 1$, then $CPA(x, 0) = x - P$, $CPA(x, 1) = x + R$, $P > R$. Similarly, $DPA(x, 0) = x - P$, $DPA(x, 1) = x + R$, $R > P$. The parameter P is called penalty, and the parameter R is called reward. Users must select the desired mean values of P and R , and the detailed organization of CPA and DPA aggregators can be obtained using appropriate software tools [17]-[19].

In some cases, propositional logic formulas must be combined with the *if-then-else* control structures [22]. That can be achieved using the selector function which compares the input degree of truth x with a threshold value T as follows:

$$b = SEL(x, T) = \begin{cases} 1, & x \geq T \\ 0 & x < T \end{cases}.$$

The selector function can be combined with the GL conjunction (C), disjunction (D), and negation (not) as shown in Fig. 2 to achieve the following general if-then-else construct:

$$z = \begin{cases} L_1(\mathbf{X}_1), & x \geq T \\ L_2(\mathbf{X}_2), & x < T \end{cases}.$$

Here $y_1 = L_1(\mathbf{X}_1) \in [0,1]$ denotes a graded propositional calculus formula based on an array of input degrees of truth \mathbf{X}_1 . Likewise, $y_2 = L_2(\mathbf{X}_2) \in [0,1]$ denotes a graded propositional calculus formula based on an array of input degrees of truth \mathbf{X}_2 . Generally, $\mathbf{X}_1 \neq \mathbf{X}_2$, but frequently we can have $\mathbf{X}_1 = \mathbf{X}_2 = \mathbf{X}$. Similar reasoning can be applied to L_1 and L_2 ; e.g., these can be the same propositional formulas that differ only in weights or only in selected inputs. The selector variables x and T can be independent inputs or selected components of arrays \mathbf{X}_1 and \mathbf{X}_2 . Obviously, the if-then-else construct provides a very high flexibility for the development of sophisticated graded propositional calculus formulas.

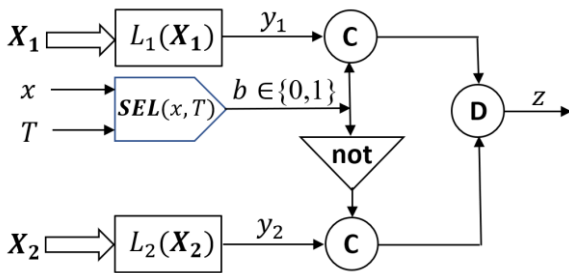


Fig 2. A general if-then-else construct implemented in GL [22]

Graded propositional calculus formulas with fixed structure and constant parameters (weights and andness/orness) are called *stationary* GL models. They are most frequently used in decision-making models that provide the overall suitability of various competitive objects/alternatives. Much less frequently we can use *nonstationary* GL models [3], where the parameters (weights, andness, orness) can be functions of input attributes. The if-then-else constructs are a special case of the nonstationary graded logic models. Of course, while the design of stationary decision models is simple and requires very modest effort, the design of nonstationary models needs significantly higher effort and a thorough justification.

V. PROFESSIONAL DECISION MAKING

A. Characteristics of professional decision making

Professional decision-making problems can be found in many areas. A detailed survey of such problems and corresponding examples can be found in [3] and [22]. Following are the main characteristics of such problems:

1. *The need for domain experts.* Graded logic decision problems cannot be successfully solved without expertise in the area of organization and functioning of evaluated objects. For example, all medical decision problems (e.g. evaluation of patient priority for organ transplantation) need the collaboration with medical doc-

tors. Evaluation teams that evaluate, compare, and select aircrafts for specific stakeholder, should include pilots. In these examples, medical doctors and pilots play the role of domain experts: they provide knowledge that is necessary to select suitability attributes, suggest importance weights, decide about the hard and soft aggregation, etc.

2. *The need for decision engineers.* Decision engineers are professional evaluators specialized in decision methods, experienced in solving decision problems, and familiar with the use of specialized software tools that are necessary for development and use of decision models. Decision engineers are central participants in evaluation teams, responsible for logic methodology, scheduling of activities and for communication with other participants in the professional decision-making team.
3. *The role of stakeholder.* In professional decision making, the stakeholder is an organization that makes decisions about selecting an object/alternative that will contribute to attaining stakeholder's goals. The stakeholder decides about development goals, provides all financing, accepts or rejects the selected best alternative, and bears the consequences of the realized decision.
4. *Organization of professional decision-making teams.* Professional decision-making teams consist of three participants: (1) stakeholder, (2) decision engineer (evaluator), and (3) domain expert. Each participant can be a single person or a group of people. In some cases, a single person can play more than one role (e.g. a decision engineer can also be a domain expert).
5. *Large number of suitability attributes.* Decision problems based on logical reasoning/evaluation can be classified according to the number of input suitability attributes. The cases below 10 inputs can be considered toy problems, popular as examples in theoretical papers. The problems below 50 inputs are frequent in individual decision making (comparison and selection of cars, homes, educational institutions, etc.). Low to medium complexity professional problems include 50 to 150 inputs. Higher complexity problems can have 150 to 600 input attributes.
6. *Precision of logic decision models.* The professional decision models should be as precise as possible. The main goal in this area is to include all suitability attributes, i.e., all components that provably affect the overall suitability of evaluated object/alternative. Equally important is to avoid considering attributes that characterize the evaluated object, but do not affect its suitability (e.g., the color of a chassis of electronic equipment is an attribute that does not affect suitability).
7. *Sensitivity and tradeoff analysis.* Before starting the use of professional decision models, it is advisable to perform a sensitivity analysis (the analysis of the impact of individual inputs on the final decision result)

and a tradeoff analysis (an analysis of compensatory properties of inputs of decision models – the capability of selected input to compensate the deficiency of another input).

8. *Reliability analysis.* The parameters of decision models (e.g., andness and importance weights) are determined by decision-making teams with limited accuracy. Reliability analysis is necessary to assess possible errors of final evaluation and selection results, and the reliability of the ranking of competitors.
9. *Optimization of evaluated objects.* For objects that have cost, stakeholders are frequently interested in solving the following optimization problems: (1) find the minimum cost necessary for achieving a specific degree of overall suitability of an evaluated object; (2) find the highest overall suitability that can be achieved with specific approved financial resources; (3) find the configuration of an evaluated object so that it yields the highest suitability obtained per invested monetary unit.
10. *Tolerance of missing data.* In some cases, the values of some suitability attributes are not available. In such cases there are two possibilities: (1) disqualify the object/alternative that has incomplete inputs, or (2) perform the decision process replacing nonexistent inputs with neutral values [20]. In most applications, the preferred method is the missingness-tolerant aggregation.
11. *The need for explainability of results.* All decisions need and can be explained in a simple verbal way. That is particularly important in professional decision making where decisions must be understood and accepted by many people in the stakeholder organization. A quantitative explainability method for evaluation decision results can be found in [21].

B. The Logic Scoring of Preference Method

Our basic concept in professional decision problems is that such problems should be solved using methods fully consistent with commonsense human logical reasoning and decision making. The method that we propose is the Logic Scoring of Preference (LSP), which is presented in detail in [3]. LSP is a human-centric decision method based on Graded Logic, organized according to observable patterns of human commonsense decision making. Consequently, it consists of the following five major steps.

1. *Identification of stakeholders and their goals.* The goal of decision making is to find the best way to satisfy goals and requirements of specific SDM. So, the initial step in the LSP method is to clearly identify the stakeholder, the purpose of evaluated objects/alternatives and the goals of evaluation and selection process. From precise identification of SDM goals and requirements it is possible to create analytic LSP decision models in correct and fully justifiable way.
2. *Development of the suitability attribute tree.* In natural commonsense decision making the number of suitability attributes is small and SDM can identify them easily

and in any order. As opposed to that, in professional decision making, it is necessary to develop a large number of attributes and that must be done in an organized and systematic way. The LSP method develops suitability attributes using a hierarchical stepwise decomposition process that starts with a single root node (overall suitability). This node is then decomposed in main components (e.g. a complex computer system can be decomposed into four main components: hardware, software, performance, and vendor support). In the next step, each component is further decomposed, creating a tree structure. At the end of decomposition process we reach components that cannot be further decomposed (e.g. the computer memory capacity is directly measurable and cannot be further decomposed). These leaves of the suitability attribute tree are suitability attributes. In a special case of a binary attribute tree with n suitability attributes, the total number of decomposable nodes is $\frac{n}{2} + \frac{n}{4} + \dots + 2 + 1 = n - 1$. So, the effort for creating a binary suitability attribute tree is proportional to $n - 1$. For non-binary trees the effort is less than the effort for the binary tree.

3. *Definition of elementary suitability attribute criteria.* For each of n suitability attributes it is necessary to create an evaluation function called the attribute criterion. E.g., if M is the memory capacity, and if $M \leq M_{min}$ is not acceptable and $M \geq M_{max}$ completely satisfies SDM's requirements, then the memory attribute suitability criterion could be the following: $x = g(M) = \min(1, \max(0, (M - M_{min}) / (M_{max} - M_{min}))$). The total effort of creating attribute criteria is proportional to n .
4. *Development of the graded logic aggregation structure.* The logic aggregation of n attribute suitability degrees follows the attribute suitability tree, going node by node from the leaves towards the root of the tree. In this process it is necessary to create $n - 1$ (or less) graded logic functions. Using these aggregation functions, the LSP method provides the graded logic criterion for computing the overall suitability $X = L(g_1(a_1), \dots, g_n(a_n))$ as a graded propositional calculus formula. The effort to complete this step is proportional to $n-1$.
5. *Computation of the overall suitability and value.* Suppose that we have $m > 1$ objects/alternatives that have costs C_1, \dots, C_m . If the approved budget is limited to C_{max} , then we assume that $C_i \leq C_{max}$, $i = 1, \dots, m$. We also assume that the overall suitability must be above the minimum threshold X_{min} , and consequently $X_i \geq X_{min}$, $i = 1, \dots, m$. All SDMs are interested in high overall suitability achieved simultaneously with the low cost. Consequently, the overall value V_i of each alternative is a hard graded conjunction of the relative suitability and the relative cost:

$$V_i = \frac{X_i}{\max(X_1, \dots, X_m)} \Delta \frac{\min(C_1, \dots, C_m)}{C_i}, \quad i = 1, \dots, m.$$

Obviously, $V_i \in [0,1]$, and such an aggregator can be a weighted geometric mean:

$$V_i = \left(\frac{X_i}{\max(X_1, \dots, X_m)} \right)^w \left(\frac{\min(C_1, \dots, C_m)}{C_i} \right)^{1-w}, \quad i = 1, \dots, m.$$

If the SDM is in situation where the suitability is more important than affordability, then $w > 1/2$. If the affordability is more important, then $w < 1/2$.

The best alternative (and the proposed decision) is the alternative/object that has the maximum value: $V^* = \max(V_1, \dots, V_m)$. In the special case $m = 1$, the single alternative is considered acceptable if $X_1 \geq X_{min}$ and $C_1 \leq C_{max}$.

The LSP method in steps 2 and 4 needs effort proportional to $n - 1$ and in step 3 the effort proportional to n . Therefore, LSP is a linear algorithm: the overall LSP effort is $O(n)$. This is a very important property because it shows that the LSP method strictly supports the human commonsense logical reasoning, but expands the applicability of this form of reasoning far beyond the natural limitations of human intuitive mental processes. That justifies the use of the LSP method in the sensitive area of professional decision making.

VI. CONCLUSION

We presented, contrasted, and confronted two approaches to the development of continuum-valued propositional logic: (1) the formal axiomatic deductive approach used in development of logic theories, and (2) the human-centric approach based on human commonsense logical reasoning with graded percepts. Our goal is to show that methods for professional decision making must be consistent with human commonsense decision making, and that human commonsense decision making is based on Graded Logic, which is the logic of natural human logical reasoning with graded percepts. Observations and applications show that the human-centric approach to logic and decision methods is a natural way to develop methods for professional decision making. In this paper we presented a condensed survey of the Graded Logic and its use in the development and use of the LSP method for professional decision making.

Graded Logic is a fully continuum-valued propositional logic: the continuum-valued variables and parameters include the graded truth, the graded importance, and the graded conjunction/disjunction. These three graded percepts are provably present in human commonsense logical reasoning. Unique properties of GL are (1) continuous transition in the whole range from the drastic conjunction to the drastic disjunction, (2) unification of complementary models of simultaneity and substitutability in a single general logic aggregator GCD, (3) selectability of annihilators, (4) selectability of idempotent or non-idempotent logic aggregators, (5) andness-directedness: visibility and adjustability of andness/orness as input parameters of the GCD function, and (6) support for stationary and nonstationary graded logic aggregators.

The LSP decision method uses all unique properties of GL to provide advanced professional decision methodology that

is fully consistent with human commonsense logical reasoning. Both GL and the LSP method have a history of successful applications, but they also offer a variety of topics for future work. These topics include LSP applications in new (particularly medical) areas, experiments with human subjects to verify and expand GL models (particularly in the areas of hyperconjunction and hyperdisjunction), the comparison of LSP with other similar methods, the development and study of a variety of nonstationary criteria, as well as the development of new decision-support software tools and their applications.

REFERENCES

- [1] Anonymous, "Logic." Wikipedia, <https://en.wikipedia.org/wiki/Logic>.
- [2] Anonymous, "Propositional calculus." Wikipedia, https://en.wikipedia.org/wiki/Propositional_calculus.
- [3] J. Dujmović, *Soft computing evaluation logic*. Hoboken, NJ 07030, USA: John Wiley & Sons, 2018.
- [4] G. Boole, G., *The Mathematical Analysis of Logic*. MacMillan, Barclay & MacMillan, Cambridge, 1847.
- [5] G. Boole, *An investigation of the laws of thought, on which are founded the mathematical theories of logic and probabilities*. Macmillan, 1853.
- [6] A. De Morgan, *Formal Logic or the Calculus of Inference Necessary and Probable*. Taylor and Walton, London 1847.
- [7] B. Steinbach, and C. Posthoff, *Logic Functions and Equations. Fundamentals and Applications using the XBOOLE-Monitor*. Springer, 2022.
- [8] D.G. Radojević, "[0,1]-valued logic: a natural generalization of Boolean logic." *Yugoslav Journal of Operations Research*, 10 (2000), No. 2, pp.185-216.
- [9] H. Rasiowa, *An algebraic approach to non-classical logics*. PWN and North-Holland, 1974.
- [10] G. Priest, *An introduction to non-classical logic*. Second edition reprinted with corrections. Cambridge University Press, 2009.
- [11] J. Dujmović, Weighted Conjunctive and Disjunctive Means and their Application in System Evaluation. *Journal of the University of Belgrade, EE Dept., Series Mathematics and Physics*, No. 483, 1974, pp. 147-158. (Available from jstor.org)
- [12] J. Dujmović, Extended Continuous Logic and the Theory of Complex Criteria. *Journal of the University of Belgrade, EE Dept., Series Mathematics and Physics*, No. 537, 1975, pp. 197-216. (Available from jstor.org)
- [13] J. Dujmović, Two Integrals Related to Means. *Journal of the University of Belgrade EE Dept., Series Mathematics and Physics*, No. 412 - No. 460, 1973, pp. 231-232. (Available from jstor.org)
- [14] J. Dujmović, Weighted Compensative Logic with Adjustable Threshold Andness and Orness. *IEEE Transactions on Fuzzy Systems*, Vol. 23, No. 2, April 2015, pp. 270-290. DOI: 10.1109/TFUZZ.2014.2312018, 2015.
- [15] J. Dujmović, Partial Absorption Function. *Journal of the University of Belgrade, EE Dept., Series Mathematics and Physics*, No. 659, 1979, pp. 156-163. (Available from jstor.org)
- [16] J. Dujmović, An Analysis of Penalty and Reward for Partial Absorption Aggregators. *Proceeding of the 2014 World Conference on Soft Computing*, pp. 126-133, Berkeley, CA, May 25-27, 2014. (Available from jstor.org)
- [17] SEAS Co., *Analysis and Synthesis of Aggregation Operators*. AnSy User Manual V.3.2. SEAS, 2010.
- [18] SEAS, Co., *AGOPcalc - Calculator for analysis and design of aggregation operators*. AGOPcalc User Manual, SEAS, 2011.
- [19] SEAS Co., *LSP.NT - LSP method for evaluation over the Internet*. LSP.NT User Manual. SEAS, 2017.
- [20] J. Dujmović, The Problem of Missing Data in LSP Aggregation. *Proceeding of the 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU, 2012*. In S. Greco et al. (Eds.): *Advances in Computational Intelligence, IPMU 2012, Part III, CCIS 299*, pp. 336-346, Springer 2012.
- [21] J. Dujmović, Interpretability and Explainability of LSP Evaluation Criteria. *Proceedings of the 2020 IEEE World Congress on Computational Intelligence*, 978-1-7281-6932-3/20, paper F-22042, July 2020.
- [22] J. Dujmović, *Graded Logic*. In preparation for Springer 2024.

Multi-Level Language Architectures as a Foundation for Advanced Enterprise Systems

Ulrich Frank

*Department of Computer Science
University of Duisburg-Essen
Germany
0000-0002-8057-1836*

Abstract—Enterprise systems are the backbone of many companies. Most operational activities are usually not feasible without them. In addition, enterprise systems may also constitute remarkable competitive advantage – or turn out to be a threat to competitiveness, depending on their quality. Enterprise systems in general, ERP systems in particular, have been around for some decades. During this time, they have undoubtedly undergone a maturing process. However, hardly any significant progress has been made regarding foundational architectures and corresponding functions. Based on an analysis of widely undisputed objectives and corresponding shortcomings of current enterprise systems, this paper presents an advanced architecture that enables the construction of *self-referential enterprise systems* (SRES). SRES promise substantial progress with respect to various essential objectives of enterprise systems. The proposed architecture is based on a multi-level language architecture. Among other things, it allows for the integration of enterprise models and corresponding software at run-time. Thus, it does not only boost reuse and adaptability, but substantially fosters user empowerment, too.

Index Terms—integration, reuse, adaptability, conceptual model, enterprise model, self-referential enterprise system, DSML

I. INTRODUCTION

TODAY’S enterprises depend on software systems. Among others, software systems are of pivotal relevance for resource management, for running business processes and for decision making. Among a plethora of specific systems, there are a few software systems that are of general relevance for a wide range of companies, e.g., systems for human resource management, for customer management, for stock management, to name a few only. The most prominent, not to say prototypical example of enterprise software are enterprise resource planning (ERP) systems. In the following, I will subsume these enterprise software systems under the umbrella term of *enterprise systems*, with specific emphasis on ERP systems.

Over the past decade, there have been various technological trends that have had an impact on the development and use of business software. Driven by the availability of non-volatile RAMs at affordable prices, the old idea of In-Memory databases [1] found its way into commercial ERP systems. It is suited to substantially increase the performance of extensive data analysis processes. Thus, it allows to integrate OLAP

and OLTP functionality in one database management system. This not only means that more up-to-date data can be used in analyses, but also makes the data warehouse system partially obsolete. Microservices represent another trend in enterprise computing. Whilst their name is misleading, they can make an important contribution to scalability and, related to the previous, to deployment of enterprise systems. Microservices may have an impact on the architecture of software systems, especially in cases where the load different parts of a system have to handle, varies. Related, but not confined to microservices is a further trend that concerns the management of enterprise systems. “Software as a Service” does not only relieve the burden on internal IT management, but is often accompanied by special billing models that may reduce overall costs – and promise better scalability. These innovations can be of great benefit and may even be a prerequisite for the realisation of certain business models. Nevertheless, they have no significant impact on the basic architecture and functionality of ERP systems.

Other trends concern development and maintenance of enterprise systems. The idea of model-driven development (MDD) has been around for some time [2]. It is based on the convincing assumption that focusing on conceptual models without the need to bother with peculiarities of implementation languages is suited to contribute to productivity of software development projects and to software quality alike. As we shall see, MDD suffers from certain pitfalls, which may have contributed to the disappointing fact that it seems not to have a substantial impact on the development and maintenance of enterprise systems. Recently, a considerable hype was generated by so called “low-code” platforms. They offer the prospect of enabling employees without specific programming knowledge to develop software. The idea is not to develop large systems, but rather to quickly create smaller systems tailored to specific needs, which are suitable for partially replacing the use of spreadsheet programs, for example. Irrespective of the fact that low-code represents a clearly exaggerated marketing trend, it can hardly be assumed that the design and functionality of business software will be influenced by it. For critical accounts of low-code platforms see [3], [4].

Our brief overview of developments that had or might have an impact on the realization and use of enterprise systems

shows that progress that concerns the principal architecture was rather limited. Even though it is interesting to ask about the reasons for this, I will refrain from analysing them here. Instead, in the following we want to explore the question of how future enterprise systems could be designed in order to offer significant advantages. To that end, I will at first look at widely undisputed objectives that enterprise systems should satisfy and identify the pivotal measure to achieve them. Against that background I present a vision of future enterprise systems, which we refer to as *self-referential enterprise system* (SRES), that goes clearly beyond the possibilities of current systems. Regarding its implementation, the vision is confronted with considerable challenges that can hardly be overcome with conventional language architectures. However, as I will show, a multi-level language architecture is suited to build and run SRES.

II. ENTERPRISE SYSTEMS: UNDISPUTED OBJECTIVES AND CHALLENGES

Various approaches to develop an idea of how to improve enterprise systems are conceivable. One could ask experienced users to report on aspects of current systems they are not satisfied with – and to express requirements future enterprise systems should fulfill. Alternatively, it would be an option to study architectures of existing systems in order to identify serious weaknesses that call for better solutions. Both approaches require considerable effort. In addition, they are accompanied by specific methodical challenges that make the success of such studies questionable. We therefore choose a different approach. Apart from specific objectives and corresponding requirements, there are various goals and related issues that should be widely agreed upon. An analysis of these objectives is not only suited to identify shortcomings of existing systems, but also to provide insights into how future systems could be designed to represent significant progress.

A. Reuse

Reuse is of pivotal relevance for the economics of enterprise systems. That does not only concern development costs, but also the effort to adapt a system to changing requirements. Especially in cases where reuse enables significant economies of scale, cost reductions can be tremendous. Reuse of software artefacts among a range of companies implies the identification of common requirements. In other words: reuse depends on *abstraction* – from specific peculiarities of certain systems onto invariant commonalities shared by a range of systems. A closer look at reuse recommends differentiating between range and productivity of reuse, also known as the power-generality trade-off [5]. The more specific a reusable artefact is, the higher is its contribution to development productivity in cases where it fits – the lower are, however, chances that it fits. On the other hand, the more generic a reusable artefact is, the higher is its potential range of reuse, hence, the achievable economies of scale. Hence, there is need to find a proper trade-off between power and generality or, certainly better, to relax this conflict of goals. Even though the idea of reuse

is especially related to software artefacts, enterprise systems should also promote the reuse of knowledge among its users. This requires to account for diverse needs and abilities of users (see Subsection II-B).

B. Accounting for Context and Perspectives

An enterprise system is not an end in itself. It is supposed to support the business. That requires accounting for relevant aspects of a company's action system, such as corporate goals, business processes, organizational structure, or decision scenarios. If the relevant context is not represented in the enterprise system, it will usually be documented separately, more or less accurate and reliable. This does not only create issues with accessibility of relevant documents, but also with their consistency. As a consequence, it is demanding to assess how well IT and business are aligned. If a changing environment demands for adapting the business or even the business model, it is required to account for both, the enterprise system and a company's action system. Again, without a representation of relevant aspects of its context, it requires additional effort to provide for conjoint change of business and IT.

Large organizations depend on separation of concerns, which translates to a variety of different professional perspectives that comprise specific goals, interests and technical languages. To provide effective support, an enterprise system should offer appropriate representations for all perspectives relevant for its users. Appropriately designed user interfaces that allow for individual adaptations are very useful in this respect. However, they are hardly sufficient to help users with gaining a deeper understanding of the system they use, the company they work for, and how their work relates to the work of others.

C. Reduction of Complexity and Need for Transparency

Enterprise systems are supposed to reduce an organization's complexity. At the same time, they contribute to a subtle increase of complexity. Often, software systems penetrate companies to a degree that many employees perceive their work through the applications they use. In other words: corporate reality is more and more constructed through enterprise systems. At the same time, to most employees the software they use remains a black box. That is not only in obvious contrast to the idea of enlightenment, which demands for a demystification of the world that surrounds us, it is also a threat to a company's competitiveness, which requires employees that are able to assess limitations and possible modifications of the systems they use. In addition, enterprise systems are part of ever growing IT infrastructures with a huge amount of different elements and dependencies between these. The resulting complexity is a clear threat to IT management and, hence, to the efficient use of IT infrastructures.

D. Integration

Integration is a prerequisite of the efficient use of enterprise systems. It requires accounting for various aspects. First, one needs to distinguish static, functional and dynamic integration.

In all cases, integration requires the affected software systems to communicate, which in turn requires common concepts, materialized, e.g., through datatypes, classes, database schemas, interface types, event types, etc. Second, similar to reuse, there is a conflict between generality and power to be accounted for. The more specific common concepts are, in other words: the more semantics they carry, the more efficient and safe communication can be, hence, the higher is the level of integration. However, the more specific concepts that enable integration are, the more systems will be excluded. This corresponds to the use of technical languages by humans. In addition to common concepts, integration of software systems also recommends the common representation of corresponding instances in order to avoid redundancy, which in turn requires common *namespaces*.

A further aspect of integration concerns organizational integration, which corresponds to IT-business alignment. Integration of this kind, too, requires common concepts shared by the two worlds. If an enterprise system requires users to know technical concepts such as file, record or module, organizational integration will be weaker than it would be with using domain-specific concepts users are familiar with. Like reuse, integration requires abstraction – on common concepts and from specific details that are peculiar to certain systems or users.

E. Adaptability

The requirements an enterprise system should satisfy may change over time. In this case, it is of crucial importance that it can be adapted with little effort and risk. At best, possible changes had been accounted for already, when a system was first designed. Ideally, this would be reflected by a software architecture that separates a presumably invariant core from possibly variable parts. If the variable parts represent monotonic extensions of the core, which is the desirable case but not trivial to achieve, changes to variable parts would not have side-effects on the core. The prerequisite of such an architecture is abstraction. Only if one succeeds in abstracting onto invariant properties of a system, these could be bundled in an invariant core. In an ideal case, changes could be performed by competent users without the need to dive into source code.

The quest for adaptability is confronted with a conflict of goals, too. It is reflected by the notions of loose and tight coupling. Loose coupling, which is favored by many as an effective measure to achieve adaptability, aims at reducing dependencies between components – in other words: it builds on generic rather than on specific interfaces – to facilitate their replacement. Abstracting onto commonalities of a range of components creates dependencies: more specific components chiefly depend on more generic ones. As long as these dependencies are invariant, they are of no harm, but of great benefit. All dependent components can be easily changed by changing the common abstraction. Fig. 1 illustrates the advantage of tight coupling in this case – and indicates the problems that arise from abstractions that turn out to be inappropriate. Related to that, there is another conflict to

account for. Adaptations of an enterprise system require some kind of language. If this language is generic, as in the case of a general-purpose programming language, a wide range of changes is possible. However, changes of this kind are very time-consuming and risky. On the other hand, a language that clearly restricts changes is likely to reduce effort and risk. Examples include approaches to configuration or domain-specific languages (DSMLs).

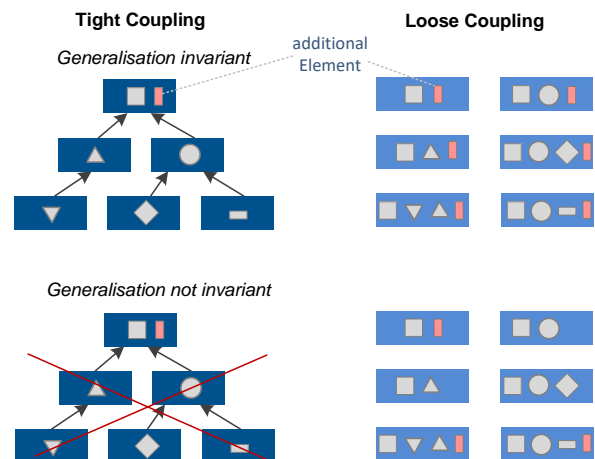


Fig. 1. Comparison of tight and loose coupling

F. Preliminary Conclusions

Our brief overview of objectives that should be widely undisputed reveals the following insights.

- Abstraction is of pivotal relevance. It is the prerequisite of reuse, integration and adaptability.
- Conceptual models are a useful instrument for developing abstractions of high quality. At best, they allow users to participate in their development and evaluation. Conceptual models could also serve as a representation that competent users could change without the need to bother with code.
- Semantics is likely to produce serious goal conflicts that require painful trade-offs. Therefore, approaches that allow mitigating these conflicts promise great benefits. Generalization/specialization is an example of how such a mitigations could work. At a higher level of generalization, a wider range of (re-) use can be expected, whereas more specific levels contribute to higher productivity, while they, at the same time, benefit from greater economies of scale through reusing higher level concepts.
- To take advantage of powerful abstractions, languages are required that provide concepts which allow for expressing these abstractions. The examples in Fig. 5 illustrate the problem. As we shall see, mainstream programming languages – and modeling languages alike – are seriously limited in this respect.

III. MULTI-LEVEL SELF-REFERENTIAL ENTERPRISE SYSTEM

Our first vision of future enterprise systems emerged some time ago. It was mainly inspired by our work on enterprise modeling. It was mainly focused on leveraging the utility of enterprise modeling tools by integrating them with enterprise software. Unfortunately, the vision suffered from serious feasibility problems. Only later, as an outcome of our research on multi-level language architectures, we were able to further refine the vision and to substantiate the design with an architecture that makes it feasible.

A. The Early Vision

The idea of enterprise modeling has been around for some time [6], [7]. It is based on the assumption that an organization's information system and its action system call for joint analysis and design in order to fully exploit the potential of IT. Therefore our work on enterprise modeling was at first focused on supporting early phases of enterprise systems' life-cycle. It resulted in a method for multi-perspective enterprise modeling (MEMO, for an overview see [8]), which includes various domain-specific modeling languages (DSMLs), e.g. for modeling corporate goals [9], [10], IT infrastructures [11], organisation structures [12], business processes [13], and decision processes [14]. These languages are integrated through a common meta-metamodel and common concepts.

Since the proper use of DSMLs as well as the analysis and management of enterprise models demand for supporting tools, we put considerable effort into the development of modeling tools [7], [15], [16]. An enterprise modeling environment such as MEMO4ADO [16] does not only allow to create the various particular models, e.g., business process models, goal models or models of the IT infrastructure. It also integrates them, thus, ensuring referential integrity of modeling elements and allowing for cross-model analysis, e.g., by allowing to navigate from a business process model to all resources that are required for its execution. Fig. 2 shows an overview of diagram types produced with MEMO4ADO and illustrates their integration through common concepts.

These benefits of a traditional environment for enterprise modeling are contrasted with serious limitations. First, models focus on the type level only. This is for a good reason. Usually, we want to intentionally fade out particular instances, since they are changing all the time. However, there are analysis scenarios where instances are important. For example, one may want to know how many instances of a certain business process types were executed within a certain month, or when a particular instance started. Other examples include the number of invoices or the invoice with the highest amount etc. To answer questions related to the instance level, one would have to use a corresponding enterprise system. If this system is not integrated with a corresponding enterprise model, it would not be possible to navigate from one system to the other – an obvious obstacle to decision making.

There is a further reason for integrating an enterprise modeling environment with an enterprise system. The development

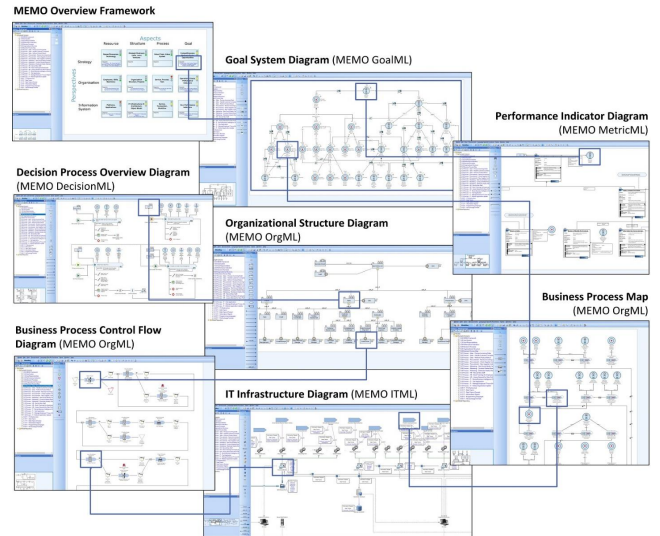


Fig. 2. Elements of MEMO4ADO

of an enterprise model requires considerable effort. Also, parts of an organizational information system are always in the state of change. Even a smoothly working IT infrastructure and well-designed business processes create the need for corresponding models in order to cope with complexity – no matter whether they are in an early or late state of their life-cycle. Hence, in order to not waste valuable resources and to enable additional benefits, we came to the conviction that it should be possible to use enterprise models during the entire life-cycle of a company and its information system, both as an instrument to support management and as a means to empower all employees by improving transparency.

This idea led quickly to the vision of integrating enterprise models – more precisely: tools for enterprise modeling – with enterprise software. We referred to this vision as “self-referential enterprise system” (SRES) [17]. An SRES results from the integration of an enterprise modeling environment and a corresponding enterprise software system. In an ideal case, developers and competent users could apply changes to parts of an enterprise model which then would become effective in the enterprise system.

To develop a demonstration of an SRES, we aimed at extending an existing enterprise modeling environment accordingly. Unfortunately, we soon had to realize that there were serious problems standing in the way of integrating the two systems. These problems were caused by principal limitations of implementation languages.

B. Challenges

These limitations create serious challenges to the design of SRES. They mainly comprise two interrelated aspects. First, mainstream programming languages do not allow for the straightforward implementation of modeling environments that represent instance level data. Second, related to that, modeling languages that are based on a MOF-like architecture

do not allow for expressing knowledge about instances. The first aspect is illustrated in Fig. 3. It shows a UML class that is conceptually located at M1. However, within a modeling tool, it is implemented as an object at M0. This is for a serious reason. A modeling tool needs to allow defining and changing the properties of a class. However, only objects have state, which can be manipulated. As a consequence, it is not possible to create instances of classes within a model editor. The only option is to generate code – resulting in two separate representations.

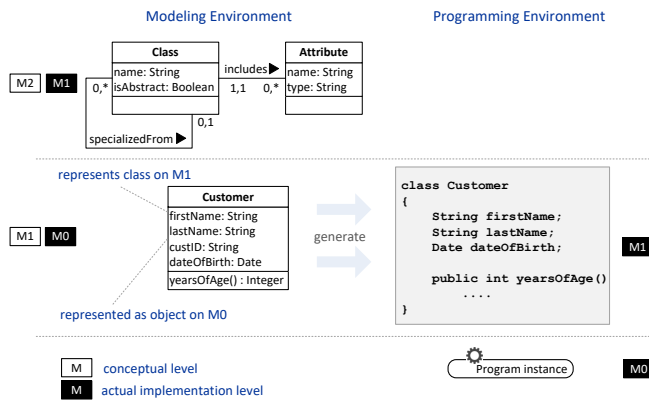


Fig. 3. The need for generating code

As a consequence of these limitations, the architecture we developed with our first conception of SRES was based on a pragmatic notion of integration. At first, concepts defined within an SRES had to be replicated in the corresponding enterprise system. That could, at best, be achieved by generating code from models. Then, both systems had to be integrated through interfaces that allow requests made in one system to be forwarded to the other system. If, e.g., a user who studies the model of an IT infrastructure within the enterprise modeling tool wants to know what instances of a certain platform type exist and where they are located, the corresponding interface should allow to send this request transparently to the enterprise system. At the same time, a process manager who is not happy with the performance of a certain business process could navigate to the corresponding process model in the enterprise modeling environment, where he might decide to change the model, which should then lead to the adaptation of the corresponding process schema in the enterprise system. Fig. 4 shows an outline of the corresponding architecture.

It is needless to say that we were not satisfied with this solution. It reflects a poor concept of integration that requires ongoing synchronization. Since new requirements often need to be implemented under time pressure, it is likely that changes are directly applied to code, which over time leads to a depreciation of the corresponding model. In addition, the second problem, namely the lack of expressiveness of modeling languages, could not be overcome with the proposed architecture. An enterprise modeling environment needs to offer DSMLs. Otherwise users would have to model goals, business processes etc. from scratch, which would not only cause unacceptable

effort, but would also be a threat to integrity. The concepts provided by a DSML serve the specification of types or classes. Fig. 5 shows fragments of two possible DSMLs and corresponding models. The concept of a printer may be part of a DSML to model IT infrastructures, whereas the concept of an activity may be offered by a DSML for modeling business processes.

Unfortunately, it is not possible to express that a particular printer has a serial number or a certain number of printed pages with the DSML, even though we know that these properties are required. Accordingly, we cannot express the knowledge that a particular business process has a start time and end time, since corresponding attributes would apply to a certain business process type, not to its instances. To express this knowledge, it would have to be added redundantly to every instance of the metaclasses **Printer** or **Activity**. However, even such a dissatisfactory approach would not allow to subsequently create particular instances within a model editor – for the reasons illustrated in Fig. 3.

A closer look at the model fragments in Fig. 5 reveals a further challenge. The specification of a metaclass like **Printer** has to be done from scratch requiring the language designer to know essential properties of a printer. Would it not be more appropriate to use an existing, more generic DSML that already includes a general concept of printer to define printer models? As we shall see such an approach would contribute to the more efficient development of DSMLs and would, at the same time, be suited to improve their quality.

IV. MULTI-LEVEL LANGUAGE ARCHITECTURES TO THE RESCUE

To cope with the limitations of the MOF architecture, we extended our previous meta modeling language with so called “intrinsic features” that allow to define features such as attributes in a meta class at M2, hence as part of a DSML, which are to be instantiated only with an instance at M0. This extension allowed, e.g., to express the fact that a particular printer has a serial number (see Fig. 5) with a DSML by characterizing the corresponding attribute as intrinsic. Unfortunately, this was little more than a Pyrrhic victory, since intrinsic features could not be expressed by common object-oriented programming languages. Furthermore, the extension was limited to M2 and there were indications already that higher levels of classification might be useful.

The back then young field of *multi-level modeling*, a term introduced more than twenty years ago by Atkinson and Kühne [18], with ancestors that go back even further, cf. [19]–[22], promised to address our needs more convincingly. However, multi-level modeling languages were not accompanied by corresponding programming languages, which we needed for our purpose. Then, about 15 year ago, a discussion at a conference dinner lead to a solution. The *XModeler*, a language engineering environment developed by Clark et al. [23], [24] proved to feature a language architecture that was suitable for a convincing implementation of SRES. Encouraged by these prospects, we started the project “Language Engineering for

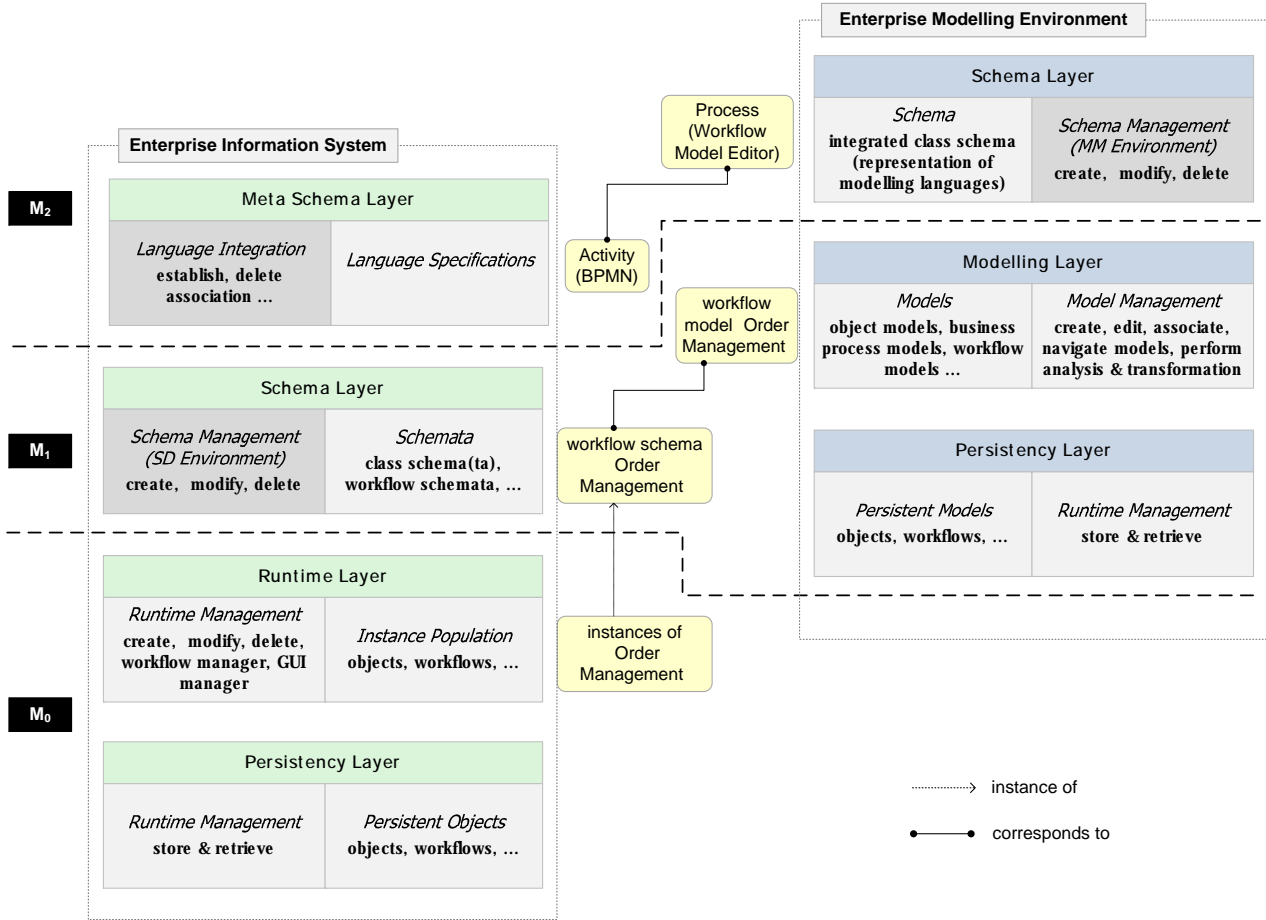


Fig. 4. Outline of early architecture of SRES

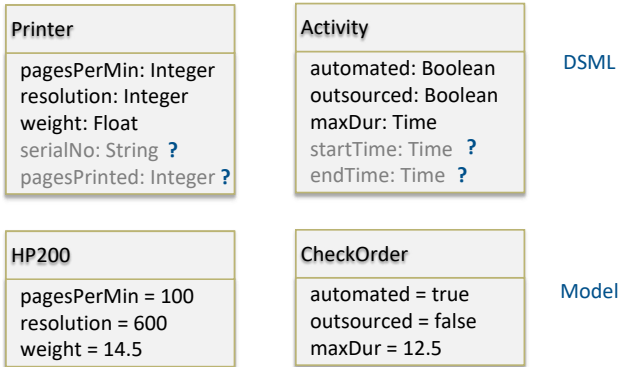


Fig. 5. Limited expressiveness of traditional languages

Multi-Level Modeling” (LE4MM, www.le4mm.org), which is still running today. For a brief history of the project see [25].

A. XModeler^{ML} and FMML^X

While the XModeler does not feature a multi-level language, its metamodel, XCore, could be easily extended to

enable essential features of a multi-level language: an arbitrary number of classes with an explicit level and deferred instantiation of properties such as attributes or operations. This extension led to the specification and implementation of FMML^X, a multi-level modeling language [26]. Different from other multi-level modeling languages such as LML [27] or M-Objects [28], FMML^X is executable, that is, it features a common representation of models and corresponding programs. The implementation of FMML^X in the XModeler led to the XModeler^{ML}. It is, together with various additional resources such as screencasts and publications, available on the project’s webpages at www.le4mm.org.

The language architecture enabled with the FMML^X allows to overcome the lack of expressiveness traditional language architectures suffer from. The small FMML^X model in Fig. 6 corresponds in part to the example in Fig. 5. It illustrates how knowledge that cannot be expressed with traditional languages can be represented without redundancy.

The class **Product** at level 3 serves the definition of properties that apply to all kinds of devices. The class **Printer** in part inherits these properties, in part instantiates them. Therefore, a specialization hierarchy would not be sufficient.

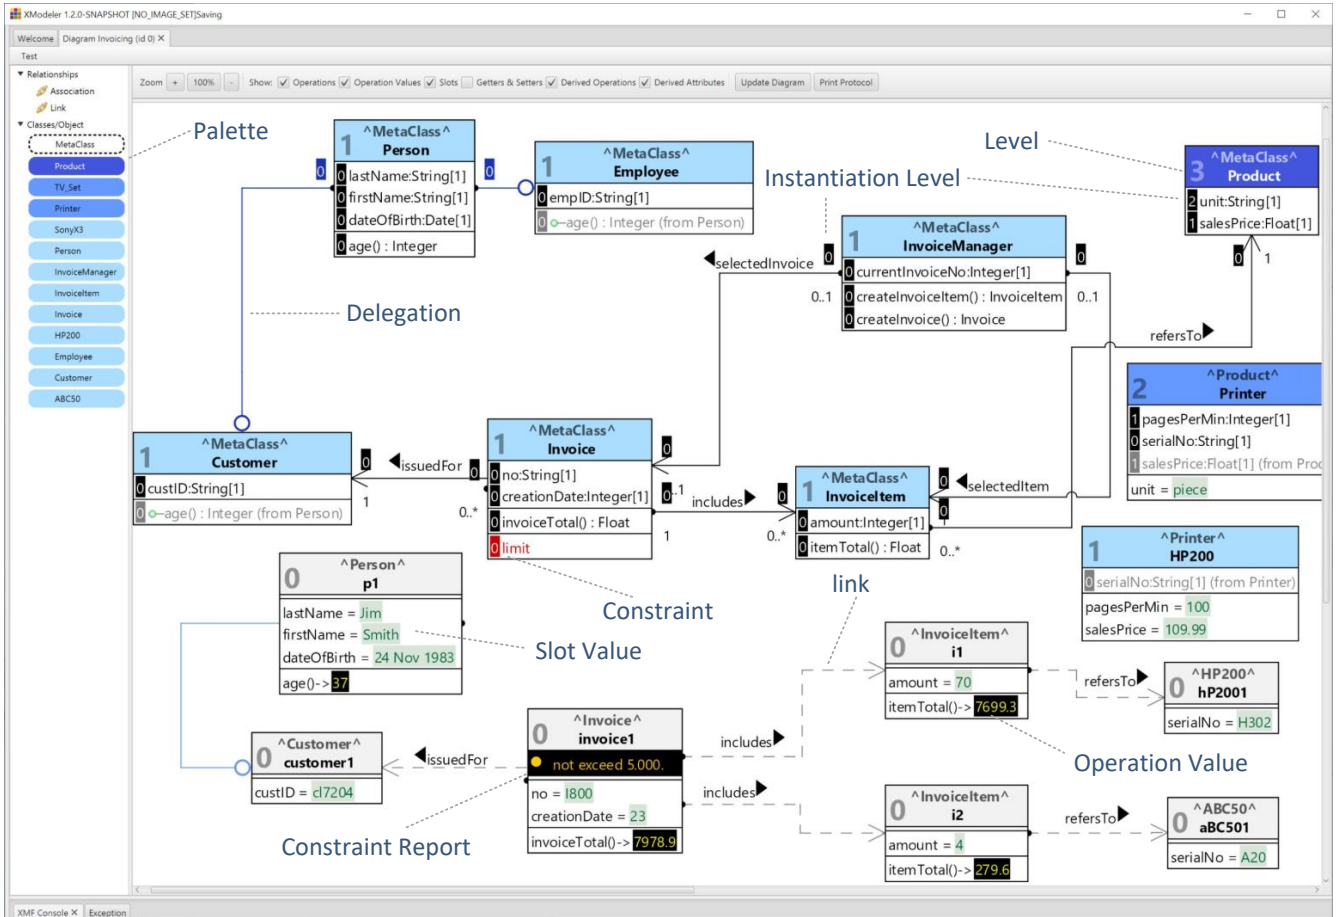


Fig. 6. Example FMML^X model created with the XModeler^{ML}

The number printed in white in a black rectangle next to an attribute or operation is to indicate the level where it is supposed to be instantiated or executed. The fact that every class is an object is, on the one hand, illustrated by slot values, such as 100 for **pagesPerMin** in the object **HP200**, which represents a printer model at level 1. On the other hand, it is shown by the values returned from the operation **invoiceTotal()** executed by instances of the classes **Invoice**, **HX500** or **Person**. In other words: an FMML^X model is executable. The objects it consists of can be displayed in a diagram editor, either with a standard or a customized notation, or by using an object browser or a customized GUI. To strengthen the integrity of the model – and of the executable program simultaneously – the FMML^X allows for adding constraints to classes, which are then immediately evaluated as soon as corresponding instances are created or changed. For example, the constraint **limit** in **Invoice** defines that the total of an invoice must not exceed 5,000, which is violated by the object **invoice1**.

The model also shows that multi-level models overcome the traditional distinction of modeling language and model. As soon as a class is created, it extends the palette and can

be used to create further instances. Furthermore, there can be links between objects at any level. A language, which is a model at a higher level, can be changed at runtime, which leads to an immediate update of the affected models. Note, however, that changes at higher levels can be challenging. Therefore, it is of crucial importance that concepts represented in a multi-level model are the more invariant the higher they are located in the hierarchy [29].

An arbitrary number of classification levels is enabled by a reflexive and recursive metamodel that specifies and implements the FMML^X. Fig. 7 shows the simplified meta model with a few selected constraints. Since **Class** inherits from **Object**, every class in the system is an object, has state and can be executed. FMML^X objects are instantiated from **MetaClass**. With their instantiation they are assigned an object of the **Level**. It allows to either define a specific level through an integer or to define a range of possible values in case the level of a class should be contingent [30]. Deferred instantiation of properties such as attributes, operations or associations is enabled through the attribute **instLevel** that serves the definition of the intended instantiation level. For a more comprehensive description of the metamodel, the related

instantiation mechanism as well as the overall architecture of the XModeler^{ML} see [31], [32].

The XModeler^{ML}, which is provided as open source, and multiple resources including screencasts, publications and example models are available on the LE4MM webpages at www.le4mm.org.

B. A Foundation for Self-Referential Enterprise Systems

A multi-level language such as the FMML^X and a corresponding language engineering and execution environment as the XModeler^{ML} provide a powerful foundation for SRES. First, they allow a common representation of models and programs. Hence, there is no need for the synchronization of two different representations. Users of an SRES can navigate from the software they use to corresponding models and meta models. If they are qualified and authorized, they may also change a model with the effect that the software they use is instantly changed, too. A multi-level model of an enterprise may be comprised of multiple DSMLs, which are defined at different levels. Since these DSMLs are all executable, they are domain-specific programming languages at the same time.

Such a language architecture supports reuse and adaptability. The knowledge represented in DSMLs can be reused. If a DSML does not fit specific needs, the language it was specified with can be used to define a new customized DSML. This allows to benefit from economies of scale supported by high level DSMLs and to benefit from the productivity provided by more specific DSMLs, thus relaxing a crucial design conflict. An SRES would then be based on a multi-level model and a corresponding runtime environment. In addition, there would be a component that serves making object models persistent and supports object retrieval. Further components would enable presentation and interaction. Since the XModeler^{ML} supports the MVC pattern, multiple views could be added to predefined diagram and browser views. Fig. 8 shows a highly simplified representation of the architecture.

C. Illustration

At an operational level, an SRES would provide GUIs similar to those known of today's ERP systems. These allow to access objects at level 0, that is, objects that represent data about particular entities or aggregations of these. In addition, an SRES would allow users to navigate to elements of the integrated enterprise model, which are typically located at level 1. These models can be presented in diagram editors featuring a standard or a customized graphical notation. In addition users could also be offered textual representations of these models. In case, users are overwhelmed by distinguishing different levels of abstraction, they could also be provided with a more traditional GUI that allows for accessing objects at different levels without the need to understand the notion of a classification level. For those users and administrators who want to understand or eventually change the DSMLs used to specify the models, an SRES would allow for accessing the full multi-level model representing an SRES. It could be represented by diagrams, within an object browser or in text

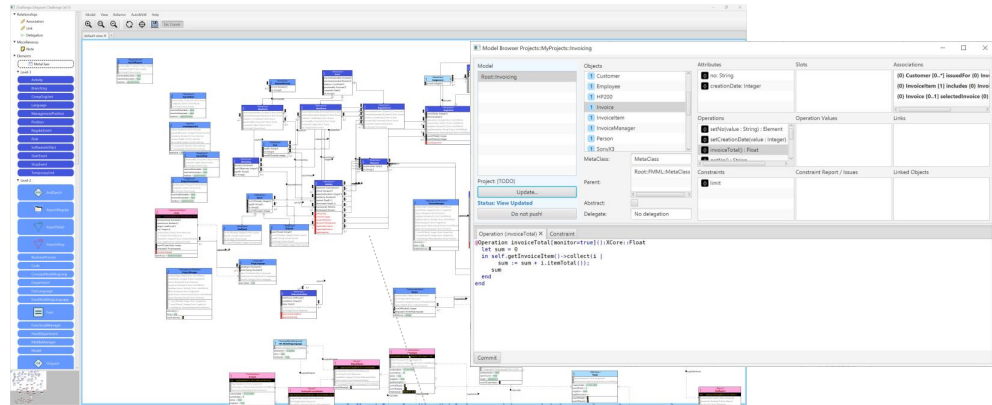
editors. The dotted edges between selected elements of the different layers are to indicate that all these representations are integrated, since they are only different views of the same multi-level system. Fig. 9 illustrates how the various levels of abstraction covered by a multi-level SRES can be presented to users.

D. Brief Evaluation

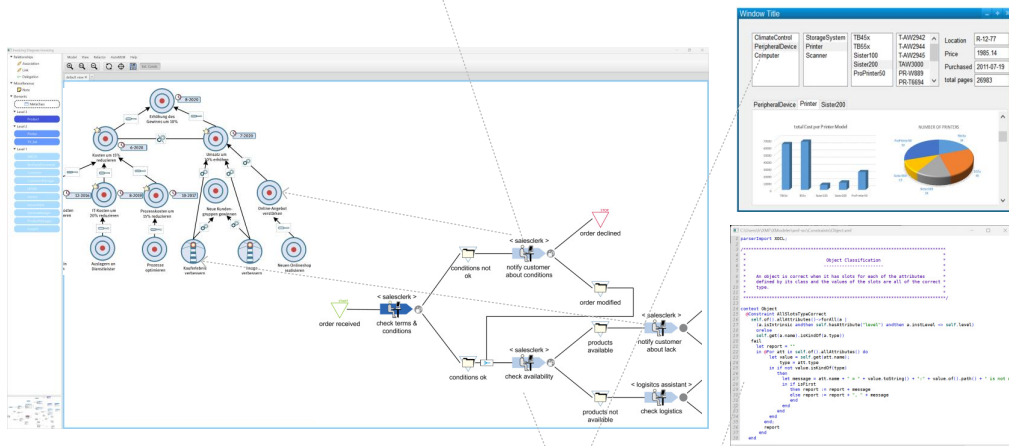
Instead of a comprehensive evaluation of multi-level architectures which can, e.g., be found in [31], I will focus on a few essential aspects only, referring to the objectives described in Section II. The integration of enterprise software with a corresponding enterprise model is obviously suited to empower users, since they have a much better chance to understand and eventually change the software they deal with. Since an SRES provides an integrated representation of company's action system and its information system, users are also supported with aligning business and IT. The complexity inherent especially to larger organizations is reduced by models that were created with DSMLs. The common representation of models and programs allows for doing without two separate representations. This does not only foster referential integrity, but also supports protection of investments into models, which otherwise are likely to be devaluated over time. Adapting an SRES to changing requirements is, at best, facilitated by applying changes at a higher level in the hierarchy only once instead of repeatedly at lower levels. In addition, a multi-level architecture also fosters reuse and, hence, economics of acquiring and managing enterprise systems. Furthermore, it also promotes cross-organizational integration of enterprise systems. Integration depends on common concepts. If, e.g., company A sends a message to company B referring to the particular printer model "HP200", communication would fail, if the software company B is using does not know a corresponding class. Within a traditional scenario, there would be no way to apply a useful interpretation of an unknown class. It would just be some class. In case of a multi-level architecture, there would be the chance to identify it as some kind of printer, if the corresponding class was known by B.

In light of these attractive prospects, it does not seem too daring to claim that multi-level architectures are suited to make enterprise software clearly more powerful. This claim leads to the obvious question why multi-level architectures have not taken off yet. There are various reasons for this unpleasant situation. First, the benefits of multi-level architectures are not easy to understand. Second, for legitimation reasons decision maker tend to opt for mature mainstream solutions. Multi-level systems are definitely not mainstream. Existing implementations of development and execution environments are restricted to academic prototypes. Third, there may be principal objections against multi-level modeling, since it may seem strange to those who are used to languages that are restricted to one classification level only. Multi-level models provide indeed features unknown of in traditional modeling and programming languages. Not only that they allow for an arbitrary number of classification levels and regard all

L2+
Multi-level model of the SRES presented as diagram or in a browser



L1
Graphical Models, specific GUIs or text editors to access models of the enterprise and of the enterprise software



L0
Traditional GUI of Enterprise System

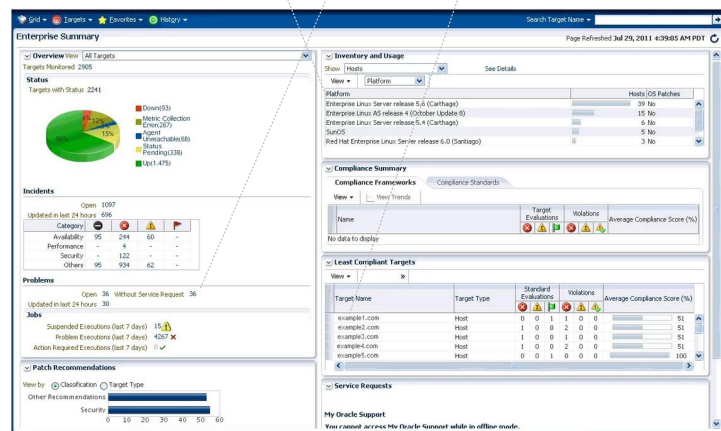


Fig. 9. SRES from user perspectives

GUI builder compared to designing a particular GUI. Reducing complexity implies increasing it first. Those who design multi-level models are confronted with remarkable complexity, especially in cases where requirements vary to a large extent. Those, however, who use an existing multi-level model and fit it to more specific needs benefit from a level of complexity that is certainly lower than that of creating a UML class diagram from scratch or dealing with representations that are used for the configuration of ERP systems.

Apart from these obstacles there are a few specific pecu-

liarities and restrictions that prevent the outlined multi-level architecture of SRES from being a silver bullet. First, the key features of a language engineering, modeling and execution environment like the XModeler^{ML}, such as an arbitrary number of classification levels and executable objects at any level are possible only through dynamic typing. Despite these obvious advantages, dynamic typing is sometimes met with reservations. Type checking happens at runtime only and, compared to languages that feature static typing, the code carries less information. It is, for example, not possible to determine the

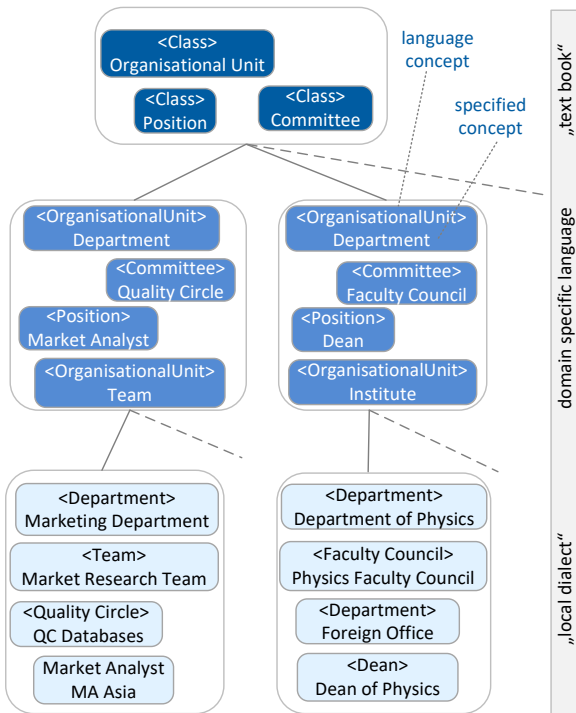


Fig. 10. Multiple levels of concepts in natural language

class of an object a message is sent to in a straightforward way. We believe that this potential advantage of statically typed languages is more than offset by the specific benefits of languages that feature *strong* dynamic typing. Languages like Smalltalk have demonstrated these benefits long ago. Looking back, it is regrettable that Smalltalk was sidelined primarily because the hardware available at the time was unable to compensate for the disadvantages. For a revealing discussion of the specific advantages of dynamic typing and its historical obstacles see this interview with Alan Kay [33]. It is needless to say that performance is not an issue anymore with today’s hardware. Second, the design of a multi-level model requires to carefully decide for a trade-off between flexibility and integrity, which is not trivial [34] (see also Subsection II-E).

Nevertheless, the realization and maintenance of multi-level language architectures is challenging. A multi-level hierarchy is extremely useful for maintaining a system as long as the dependencies reflected by the hierarchy – lower level objects depend existentially on higher level classes – are invariant over time. Therefore, the design of multi-level models requires a high level of expertise and great care. Otherwise, the advantage turns into a serious problem. A further aspect is of utmost relevance with respect to the power of multi-level language architectures. So far, multi-level models are widely restricted to static and, to a lesser degree, functional abstractions. It is much harder to define multi-level semantics for dynamic abstractions, e.g. process models. There are important reasons for this, such as the fact that specialization of process types cannot be defined as monotonic extension of a super process

type. As a consequence, the substitutability constraint cannot be satisfied – with serious implications for the maintenance of larger process landscapes. There are a few contributions to multi-level process modeling, e.g., [35], [36], [37], [38], but their main focus is on abstracting on static or functional aspects of processes. While the missing support for dynamic abstractions does not invalidate the benefits gained from multi-level static abstractions, it clearly emphasized the need for corresponding research.

V. CONCLUSIONS AND FUTURE RESEARCH

While ERP systems are of pivotal relevance for many companies’ competitiveness, only little research on future enterprise software systems happens in academia. At the same time, progress of commercial systems remains modest, at least with respect to principal functionality. Enterprise modeling, on the other hand, has seen more than two decades of research in academia, but only little adoption in business. Nevertheless, the potential benefits of enterprise models are widely undisputed. The presented architecture of SRES is suited to promote the utility of enterprise models and, at the same time, improve the power of enterprise software. While the implementation of SRES is widely impossible with prevalent language technologies, multi-level languages and corresponding development and execution environments provide a solid foundation for that purpose. In addition to enabling SRES, multi-level language architectures also allow for enriching other types of software with additional abstraction.

Our future research is primarily characterized by two directions. On the one hand, we will continue to work on concepts that allow for multi-level dynamic abstractions. In doing so, we are thinking about developing a relaxed concept of specialization. On the other hand, our work aims to simplify the transition to multi-level modeling by supporting the step-by-step enrichment of a UML editor with further concepts up to the XModeler^{ML}. A first prototype of this UML editor, called “UML-MX” is available on the project’s webpage at <https://www.wi-inf.uni-due.de/LE4MM/uml-pp/>. Among other things, it allows the instantiation and execution of objects from a UML class diagram within the model editor.

REFERENCES

- [1] M. H. Eich, “Mars: The Design of a Main Memory Database Machine,” in *Database Machines and Knowledge Base Machines*, ser. The Kluwer International Series in Engineering and Computer Science, Parallel Processing and Fifth Generation Computing, M. Kitsuregawa and H. Tanaka, Eds. Boston, MA: Springer, 1988, vol. 43, pp. 325–338.
- [2] R. B. France and B. Rumpe, “Model-driven Development of Complex Software: A Research Roadmap,” in *Workshop on the Future of Software Engineering (FOSE ’07)*, L. C. Briand and A. L. Wolf, Eds. IEEE CS Press, 2007, pp. 37–54.
- [3] A. C. Bock and U. Frank, “Low-Code Platform,” *Business & Information Systems Engineering*, vol. 63, no. 6, pp. 733–740, 2021.
- [4] J. Cabot, “Positioning of the Low-Code Movement within the Field of Model-Driven Engineering,” in *Proceedings of the 23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems*. IEEE, 2020, pp. 535–538.

- [5] A. C. Bock, "The Power/Generality Trade-Off in Decision and Problem Modeling: Theoretical Background and Multi-level Modeling as a Resolution," in *Enterprise, Business-Process and Information Systems Modeling*, ser. Lecture Notes in Business Information Processing, J. Gulden, I. Reinhartz-Berger, R. Schmidt, S. Guerreiro, W. Guédria, and P. Bera, Eds. Cham: Springer International Publishing, 2018, vol. 318, pp. 213–228.
- [6] J. A. Zachman, "A Framework for Information Systems Architecture," *IBM Systems Journal*, vol. 26, no. 3, pp. 276–292, 1987.
- [7] U. Frank, *Multiperspektivische Unternehmensmodellierung: Theoretischer Hintergrund und Entwurf einer objektorientierten Entwicklungsumgebung*. München: Oldenbourg, 1994.
- [8] —, "Multi-Perspective Enterprise Modeling: Foundational Concepts, Prospects and Future Research Challenges," *Software and Systems Modeling*, vol. 13, no. 3, pp. 941–962, 2014.
- [9] S. Overbeek, U. Frank, and C. A. Köhling, "A Language for Multi-Perspective Goal Modelling: Challenges, Requirements and Solutions," *Computer Standards & Interfaces*, vol. 38, pp. 1–16, 2015.
- [10] A. Bock and U. Frank, "MEMO GoalML: A Context-Enriched Modeling Language to Support Reflective Organizational Goal Planning and Decision Processes," in *Conceptual Modeling: 35th International Conference, ER 2016*, I. Comyn-Wattiau, K. Tanaka, I.-Y. Song, S. Yamamoto, and M. Saeki, Eds. Cham: Springer, 2016, pp. 515–529.
- [11] U. Frank, M. Kaczmarek-Heß, and S. D. Kinderen, "IT Infrastructure Modeling Language (ITML): A DSML for Supporting IT Management. ICB Report No. 71, University of Duisburg-Essen."
- [12] U. Frank, "MEMO Organisation Modelling Language (1): Focus on Organisational Structure."
- [13] —, "MEMO Organisation Modelling Language (2): Focus on Business Processes. ICB Research Report No. 49., University of Duisburg-Essen." 2011.
- [14] Alexander Bock, "Beyond Narrow Decision Models: Toward Integrative Models of Organizational Decision Processes," in *Proceedings of the 17th IEEE Conference on Business Informatics (CBI 2015)*, D. Aveiro, U. Frank, K. J. Lin, and J. Tribolet, Eds., Lisbon, 2015.
- [15] J. Gulden and U. Frank, "MEMOCenterNG – A Full-Featured Modeling Environment for Organisation Modeling and Model-Driven Software Development," in *Proceedings of the 2nd International Workshop on Future Trends of Model-Driven Development (FTMDD 2010)*, 2010.
- [16] A. Bock, U. Frank, and M. Kaczmarek-Heß, "MEMO4ADO: A Comprehensive Environment for Multi-Perspective Enterprise Modeling," in *Proceedings of the Modellierung 2022 Satellite Events*, J. Michael, J. Pfeiffer, and A. Wortmann, Eds. Bonn: GI, 2022, pp. 245–255.
- [17] U. Frank and S. Strecker, "Beyond ERP Systems: An Outline of Self-Referential Enterprise Systems: Requirements, Conceptual Foundation and Design Options. ICB Research Report No. 31. University of Duisburg-Essen," Essen.
- [18] C. Atkinson and T. Kühne, "The Essence of Multilevel Metamodeling," in *UML 2001 - The Unified Modeling Language. Modeling Languages, Concepts, and Tools*, ser. Lecture Notes in Computer Science, M. Gorgolla and C. Kobryn, Eds. Berlin and London, New York: Springer, 2001, pp. 19–33.
- [19] J. J. Odell, "Power Types," *Journal of Object-Oriented Programming*, vol. 7, no. 2, pp. 8–12, 1994.
- [20] R. C. Goldstein and V. C. Storey, "Materialization," *IEEE Transactions on Knowledge and Data Engineering*, vol. 6, no. 5, pp. 835–842, 1994.
- [21] A. Pirotte, E. Zimányi, D. Massart, and T. Yakusheva, "Materialization: A Powerful and Ubiquitous Abstraction Pattern," in *Proceedings of the 20th International Conference on Very Large Data Bases*, ser. VLDB '94, J. B. Bocca, M. Jarke, and C. Zaniolo, Eds. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc, 1994, pp. 630–641.
- [22] M. Jarke, S. Eherer, R. Gellersdörfer, M. Jeusfeld, and M. Staudt, "ConceptBase – A Deductive Object Base for Meta Data Management," *Journal of Intelligent Information Systems*, vol. 4, no. 2, pp. 167–192, 1995.
- [23] T. Clark, P. Sammut, and J. S. Willans, "Super-Languages: Developing Languages and Applications with XMF (2nd ed.)," *CoRR*, 2015. [Online]. Available: <http://arxiv.org/abs/1506.03363>
- [24] T. Clark, P. Sammut, and J. Willans, *Applied Metamodeling: A Foundation for Language Driven Development*, 2nd ed. Ceteva, 2008.
- [25] U. Frank and T. Clark, "Language Engineering for Multi-Level Modeling (LE4MM): A Long-Term Project to Promote the Integrated Development of Languages, Models and Code," in *Proceedings of the Research Projects Exhibition at the 35th International Conference on Advanced Information Systems Engineering (CAiSE 2023)*, ser. CEUR, J. Font, L. Arcega, J.-F. Reyes-Román, and G. Giachetti, Eds., 2023, pp. 97–104.
- [26] U. Frank, "The Flexible Multi-Level Modelling and Execution Language FMML^X. ICB Research Report No. 66. University of Duisburg-Essen," Essen.
- [27] C. Atkinson and R. Gerbig, "Flexible deep modeling with melanee," in *Modellierung 2016, 2.-4. März 2016, Karlsruhe - Workshopband*, ser. Modellierung 2016, S. B. U. Reimer, Ed., vol. 255. Bonn: Gesellschaft für Informatik, 2016, pp. 117–122. [Online]. Available: <http://subs.emis.de/LNI/Proceedings/Proceedings255/117.pdf>
- [28] B. Neumayr, K. Grün, and M. Schrefl, "Multi-level domain modeling with m-objects and m-relationships," in *Proceedings of the 6th Asia-Pacific Conference on Conceptual Modeling (APCCM)*, S. Link and M. Kirchberg, Eds. Wellington: Australian Computer Society, 2009, pp. 107–116.
- [29] U. Frank, "Prolegomena of a Multi-Level Modeling Method Illustrated with the FMML^X," in *Proceedings of the 24th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings*. IEEE, 2021.
- [30] U. Frank and D. Töpel, "Contingent Level Classes: Motivation, Conceptualization, Modeling Guidelines, and Implications for Model Management," in *Proceedings of the 23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings*, E. Guerra and L. Iovino, Eds. New York, NY, USA: ACM, 2020, pp. 622–631.
- [31] U. Frank, "Multi-level Modeling: Cornerstones of a Rationale," *Software and Systems Modeling*, vol. 21, no. 1, pp. 451–480, 2022.
- [32] T. Clark and J. Willans, "Software Language Engineering with XMF and XModeler," in *Computational linguistics*, I. R. Management Association, Ed. Hershey, Pennsylvania (701 E. Chocolate Avenue, Hershey, Pa., 17033, USA): IGI Global, 2014, pp. 866–896.
- [33] S. Feldman, "A Conversation with Alan Kay," *Queue*, vol. 2, no. 9, pp. 20–30, 2004.
- [34] U. Frank and T. Clark, "Peculiarities of Language Engineering in Multi-Level Environments or: Design by Elimination," in *Kühne (Ed.) 2022 – Proceedings of the 25th International*, pp. 424–433.
- [35] B. Neumayr, C. G. Schuetz, and M. Schrefl, "Dual deep modeling of business processes: 7:1-31 pages / enterprise modelling and information systems architectures (emisaj), vol. 17 (2022)," 2022.
- [36] A. Lange and C. Atkinson, "Multi-level Modeling with LML. A Contribution to the MULTI Process Challenge," *Enterprise Modelling and Information Systems Architectures (EMISAJ)*, vol. 17, pp. 1–36, 2022.
- [37] M. A. Jeusfeld, "Evaluating DeepTelos for ConceptBase: A Contribution to the MULTI Process Challenge," *Enterprise Modelling and Information Systems Architectures (EMISAJ)*, vol. 17, 2022.
- [38] U. Frank and T. Clark, "Multi-Level Design of Process-Oriented Enterprise Information Systems," *Enterprise Modeling and Information Systems Engineering (EMISAJ)*, vol. 10, pp. 1–50, 2022.

Machine Learning in Energy and Thermal-aware Resource Management of Cloud Data Centers: A Taxonomy and Future Directions

Shashikant Ilager^{1,2} and Rajkumar Buyya¹

¹Cloud Computing and Distributed Systems (CLOUDS) Lab
School of Computing and Information Systems
University of Melbourne, Australia

²Institute of Information Systems Engineering
TU Wien, Austria

Abstract—Cloud data centres (CDCs) are the backbone infrastructures of modern digital society, but they also consume huge amounts of energy and generate heat. To manage CDC resources efficiently, we must consider the complex interactions between diverse workloads and data centre components. However, most existing resource management systems rely on simple and static rules that fail to capture these complex interactions. Therefore, we require new data-driven Machine learning-based resource management approaches that can efficiently capture the interdependencies between parameters and guide resource management systems. This review describes the in-depth analysis of the existing resource management approaches in CDCs for energy and thermal efficiency. It mainly focuses on learning-based resource management systems in data centres and also identifies the need for integrated computing and cooling systems management. A taxonomy on energy and thermal efficient resource management in data centres is proposed. Furthermore, based on this taxonomy, existing resource management approaches from server level, data centre level, and cooling system level are discussed. Finally, key future research directions for sustainable Cloud computing services are proposed.

Index Terms—Cloud Computing, Energy Efficiency, Thermal-aware Workload Management, Sustainable Computing, Machine Learning

I. INTRODUCTION

CLOUD computing has changed the way computing services are delivered to end-users by providing flexible and on-demand access to resources with a pay-as-you-go model [1], [2]. Cloud computing follows the principle of providing computing resources as utility services (e.g., water and electricity). This unique and flexible service delivery model ensures that individuals and businesses can easily access the required computing services. By default, Cloud workloads require continuous, always-on, and 24×7 access to its deployed services. For instance, the Google search engine is expected to achieve an almost 100% availability rate [3]. Similarly, Amazon AWS witnesses thousands of Elastic Compute

(EC2) instances created [4] in a day through their automated APIs, thus requiring CDCs to support such critical demand. According to Gartner, by 2022, 60% of organisations will use an external Cloud service provider [5], and by 2024, Cloud computing alone will account for 14.2% of total global IT spending [6].

Cloud computing services are broadly categorised into three types. First, the Infrastructure as a Service (IaaS) model offers computing, storage, and networking resources either in virtual or physical form. Second, the Platform as a Service (PaaS) model offers tools for rapid application development and deployment, such as middleware platforms, Application Programming Interfaces (APIs), and Software Development Kits (SDKs). Third, the Software as a Service (SaaS) model offers direct access to application software to the users, and the software is developed and managed by service providers completely.

All of these service paradigms rely on the data centres to deliver the resources required for the applications and users seamlessly. Cloud Data Centres (CDCs) are massive network-based infrastructures managed in runtime by Resource Management Systems (RMS). Fig. 1 shows an abstract view of data centre infrastructure and its resource management system. The DCs host thousands of servers, networking equipment, and cooling systems. Servers and networking equipment provide the required computational resources for cloud users, and the cooling system helps to remove the heat generated by the computing resources. An RMS in the data centre is a software platform that manages different subsystems in the data centre through various tasks, such as resource monitoring, provisioning, workload scheduling, and placement. It also controls power and cooling management knobs. Some public cloud service providers build their own in-house RMS, while many private and public clouds use open-source systems such as OpenStack¹.

Note: This work is done when first author was working at University of Melbourne, Australia

¹<https://www.openstack.org>

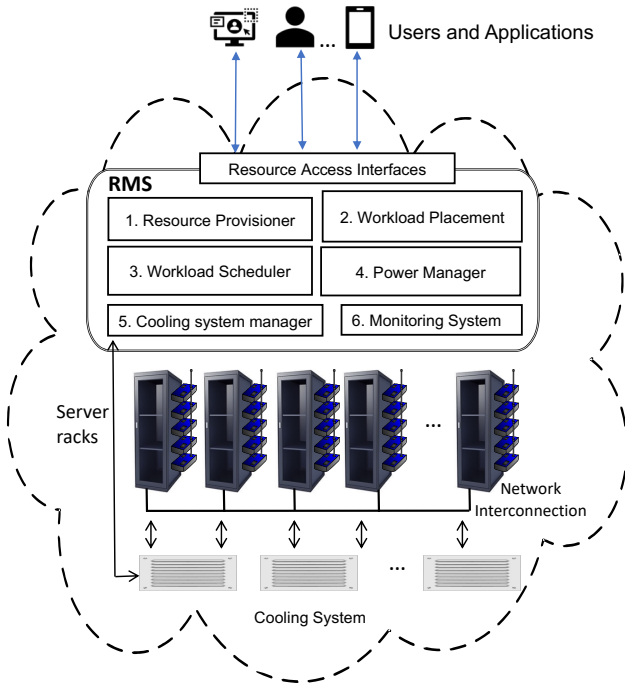


Fig. 1: An abstract view of a CDC. It shows high level view of an RMS tasks required to manage the CDCs resources and user workloads (adapted from [7]).

To meet the demand for Cloud services, major Cloud services providers such as Amazon AWS², Microsoft Azure³, and Google Cloud⁴ are deploying many hyper-scale data centres in multiple regions worldwide. There are over 8 million data centres globally, ranging from private small-scale to hypers-scale DCs, and they are growing at 12% annually [8]. As they grow in number and size, they consume more energy and face massive energy challenges. CDCs consume an estimated 2% of global electricity generated [9] and rely on fossil-fuel-based or brown energy sources that emit 43 million tons of CO₂ per year and increase at 11% annually [10], leaving high carbon footprints. Therefore, improving the energy efficiency of Cloud data centres is vital for sustainable and cost-effective Cloud computing. DCs' tremendous growth has introduced massive energy challenges. If not addressed, data centres may consume up to 8000 terawatts of power by 2030 in the worst case. However, if best practices are adopted across the Cloud computing stack, this energy consumption can be reduced to around 1200 terawatts [11] (see Fig. 2). To achieve this best-case scenario, energy-efficient practices are needed in various levels of data centre resource management platforms (such as optimised use of computing and cooling resources). Hence, addressing this energy problem and achieving sustainability, both environmentally and economically, is crucial.

²<https://aws.amazon.com/>

³<https://azure.microsoft.com/>

⁴<https://cloud.google.com/>




Fig. 2: Estimation of Data Centre Energy Consumption by 2030 [11]

A data centre is a complex cyber-physical system (CPS) that consists of thousands of rack-mounted physical servers, networking equipment, sensors, cooling systems, and other facility-related subsystems. It consumes up to 30-40 kW per rack and generates a lot of heat, posing a serious challenge for efficient and reliable resource and energy management. In particular, the main power consuming subsystems in CDCs are the computing and cooling systems, which together account for 85% of total energy consumption in a data centre, with each of them significantly contributing [12] to the total power consumption. Therefore, there is an essential requirement for integrated Energy and Thermal-aware Resource Management.

Traditionally, cooling system management and computing system management are done separately by the facility management team and the IT administrator, respectively. However, optimising one system may have a negative impact on the other system. For example, increasing resource utilisation in computing may create hotspots, and thus increase cooling energy costs. Therefore, managing these subsystems independently may result in energy inefficiencies in the data centres even if they are individually optimised for energy efficiency. The advancement in IoT and smart systems [13] has enabled many mechanical systems associated with cooling to be controlled or configured through software systems [14]–[16]. Hence, it is crucial to apply resource management techniques holistically to optimise both computing and cooling systems and avoid conflicting trade-offs between these two subsystems.

Resource management in data centres is highly challenging due to the complex interactions between subsystems and the heterogeneous characteristics of workloads. Manual fine-tuning of the controllable parameters by resource management systems is infeasible. For example, “Just 10 pieces of equipment, each with 10 settings, would have 10 to the 10th power, or 10 billion possible configurations, a set of possibilities far beyond the ability of anyone to test for real” [17], [18]. Moreover, these large-scale systems often have nonlinear relationships between their parameters. However, optimising data centre operation requires adjusting

the hundreds of parameters in different subsystems where heuristics or static solutions are ineffective.

Therefore, to cope with the complexity of data centre infrastructures and the dynamic nature of cloud workloads, Machine Learning (ML)-based resource management methods are vital. In parallel, integrated resource management of the computing and cooling systems is necessary to balance the trade-offs between these two subsystems and achieve significant energy efficiency in CDCs [7]. There have been many efforts in this direction using ML for systems focusing on optimising different computing systems [19]. For instance, ML-centric Cloud [20], developed Resource Control (RC), a general ML and prediction serving system that provides insight into the Azure compute fabric resource manager’s workload and infrastructure. Similarly, Google has used ML models to optimise the efficiency of its data centres by adjusting the different knobs of the cooling system, thus saving a significant amount of energy [21]. These use cases demonstrate the feasibility and benefits of learning-based solutions in different aspects of resource management in clouds. Moreover, even a 1% improvement in data centre efficiency can save millions of dollars per year and reduce the carbon footprint [22].

The rest of the paper is organised as follows: Section II provides overview of ML-based RMS in CDCs. Section III and Section IV review the existing methods for energy and thermal management in data centres based on the taxonomy, respectively. Section V explains the integrated resource management solutions for energy and thermal efficiency. Section IV-C describes different cooling systems in a data centre, including air and liquid cooling systems. Section VI outlines the future research directions. Finally, Section VII concludes the paper.

II. BACKGROUND: ML-BASED RESOURCE MANAGEMENT SYSTEMS IN CDCS

Machine learning (ML) is naturally used in Computer Vision (CV) and Natural Language Processing (NLP) problems due to its ability to identify patterns from the complex input data. ML algorithms are classified into supervised and unsupervised learning, depending on the input data preparation and training methods. ML methods itself can be broadly used for numerical prediction- *regression models*, and for categorization based on class labels,- *classification modes*, as well for developing advanced control systems- *Reinforcement Learning (RL) controllers*.

As data center complexities increase, ML algorithms are required to perform a variety of RMS tasks. For instance, as illustrated in Figure 3, the left side of the Figure 3 indicates the high level RMS Tasks in a CDC (also see Figure 1). Please note that, these tasks indicate primary functionality of an CDCs middleware system; there could be other tasks based on data center and workload requirements. Similarly, the right side of the Figure indicates list of all possible ML tasks an RMS could require in its decision making process. For example, the *Resource Provisioner* can invoke resource estimation models to predict the required amount of computing

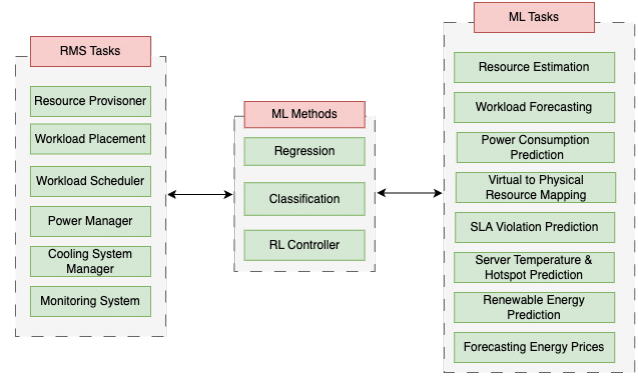


Fig. 3: List of RMS Tasks and ML Tasks and ML Methods.

resources for a workload. The RMS can also be guided by power consumption and SLA violation prediction models, depending on the optimization objectives.

Much advanced ML applications such as RL can be used to develop controller systems. An RMS can be modeled as a decision engine with a list of actions designed to satisfy specific objectives. Such approaches are increasingly being used in CDCs for cooling system knob configuration, scheduling, and power management systems. Although Figure 3 provides an overview of how ML can be leveraged to develop highly optimised RMS systems for today’s complex CDCs, it is not exhaustive.

III. TAXONOMY OF ENERGY MANAGEMENT IN CDCS

Many researchers have focused on increasing the energy efficiency of data centres with various resource management techniques. These techniques cover an individual server to geo-distributed data centres. Taxonomy on the data centre’s energy management solutions is presented in Fig. 4. We categorise these solutions into two broad categories, i.e., single server level and data centre level solutions. Accordingly, we identify the essential techniques used in these two categories and briefly review their methods.

A. Server Level

In a computing server, the CPU predominantly consumes a significant amount of energy. Modern rack-mounted data centre servers consume more than 1000 watts of power. Hence, managing this high power consumption is a challenging task. This server-level power management has been mostly left to the operating system and its device drivers that communicate with underlying hardware signals and manage the server power. Server-level power management can be broadly categorized into two levels, static and dynamic power management. Static power management deals with minimising leakage power, while dynamic power management deals with regulating active runtime power based on utilization level.

1) *Static Power Management*: The silicon chip has static power consumption, which is independent of the usage level. Static power mainly accounts for the leakage of current inside active circuits. To some extent, static power consumption

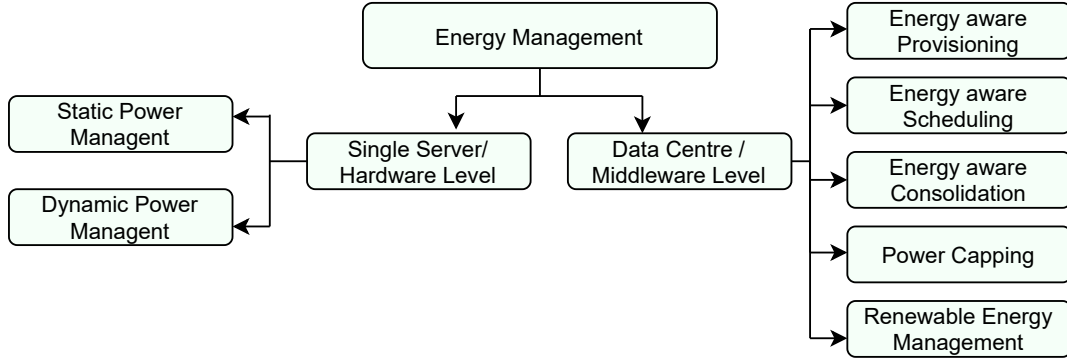


Fig. 4: Taxonomy of Energy Management in Cloud Data Centres

is unavoidable; however, it can be minimized with better design and processes. There are many solutions from a lower level from circuit level and architectural techniques [23]. The general approach in managing leakage is with different sleep states of CPUs when the system is idle. For instance, Intel X86 architecture has (C0-C4) sleep states indicating C0 is an active state, while C4 is a deep sleep state where most of the CPUs' components are turned off to avoid static power consumption. This processor's sleep state management is usually done in reactive manner at Operating System's (OS) kernel level. If a processor core is idle predefined time interval, a kernel governor changes the sleep state.

Application of ML: However, existing reactive static power management approaches could be vastly improved using the ML-based solutions with proactive strategies. For instance, Chung et al. [24] proposed a power management technique for an arbitrary number of sleep states, which turns off idle processors based on idle period clustering and adaptive learning trees. Instead of predefined interval, they estimate the adaptive intervals based on recent history, saving the energy consumption. Similarly, Lu et al. [25] introduce RAMZzz, a memory system design which is based on rank-aware energy-saving optimizations, in memory systems. It groups pages with similar access patterns into the same rank, allowing for dynamic page migrations to optimize access locality and employs adaptive state demotions with a prediction model to increase the energy efficiency. These studies indicates that ML has been widely getting used at very low level management of hardware devices.

2) *Dynamic Power Management (DPM):* A large part of silicon chip-based computing elements, either in CPU or GPUs spend on dynamic power. Dynamic power represents runtime energy based on workload utilisation level. CPUs operate at different frequencies to regulate the dynamic power. If the operating frequency of a CPU is the highest, then its dynamic power consumption will also be higher. The frequency is regulated based on utilisation level and workload requirements to increase their speedup. Dynamic Voltage Frequency Scaling (DVFS) is a popular technique to regulate the dynamic power in modern systems [26]. The dynamic power can be defined

as below:

$$P_{dynamic} \propto V^2 F \quad (1)$$

In Equation 1, F is the frequency, and V is the supply voltage to the processor. Based on the frequency, the voltage is regulated, and some frequency ranges usually have a similar. If a CPU should be at its highest speed or frequency should be set to a higher level, thus consuming more power. The operating system scales frequency based on its workload and application demands in runtime.

The DVFS-based optimizations are employed using application metrics, VM-level metrics, or even data-center level utilization metrics [27], [28]. A few studies proposed DVFS techniques at the data center level. These solutions include DVFS-aware VM scheduling, consolidation [27], [28], placement of application based on DVFS capabilities [29], and data centre level task scheduling by synchronizing the frequency scaling among multiple machines [30]. All of these works use heuristic based solutions.

Application of ML: ML-based techniques widely used in DPM optimisations. The Authors in [31] proposed ML-based CPU and GPU DVFS regulator for compute-heavy mobile gaming applications that coordinates and scale frequencies with performance and energy improvements. Similarly, in one of our recent study, we explored how ML-algorithms can help us to dynamically configure GPUs clock frequency based on workload requirements such as deadline [32]. Here, we used popular GPU benchmarks Rodinia and Polybench and collected profiled data which includes hardware level counters. This collected data is further used develop ML regression models to estimate power consumption and execution time across configurable memory and streaming processors frequencies. These models are further used to guide a scheduling algorithm to execute application within predefined deadline with minimal energy consumption.

B. Data Centre Level

A significant amount of energy efficiency can be achieved when data centre-level platforms incorporate energy-efficient resource management policies. Distributed data centre applications span hundreds of machines in geo-distributed data

centres; hence, providing energy efficiency holistically across data centre resources and applications is more feasible and yields better results. In this section, we discuss important techniques for data centre-level energy-efficient solutions.

1) *Energy-aware Provisioning*: Cloud data centres offer computing resources in terms of Virtual Machines (VMs) or containers. Allocating the required amount of resources for the application need is vital to satisfy the SLAs. However, over-provisioning of resources may yield higher energy consumption and monetary cost to the users, while under-provisioning will potentially violate the SLAs. Many researchers have proposed energy-aware resource provisioning techniques. Authors in [33] investigated energy-aware resource allocation for scientific applications. The proposed system EnReal leverages the dynamic deployment of VMs for energy efficiency. Similarly, Li et al. [34] proposed an iterative algorithm for energy-efficient VM provisioning for application tasks. Beloglazov et al. [35] propose various heuristic algorithms for resource allocation policies for VMs defining architectural principles.

Application of ML: Mehriar et al. [36] offered clustering and prediction-based techniques; they used K-means for workload clustering and stochastic Wiener filter to estimate the workload level of each category and accordingly allocate resources for energy efficiency. Recently Microsoft has proposed Resource Control (RC) [20], where they trained ML models to output predictions like VM lifetime, CPU utilisation, and maximum deployment of VMs. These predictions use various resource management problems for better decision-making, including resource provisioning with the right container size for applications. With increasing availability of data in cloud platforms in regard to user workload behaviours, and usage patterns, ML will be key technique to estimate right amount of computing powers required for user requests.

2) *Energy-aware Scheduling*: Scheduling is a fundamental and essential task of a resource management system in Cloud data centres. It addresses the following question, given an application or set of VMs (considering the application runs inside these isolated VMs), when and where to place these VMs/applications among available physical machines? This decision depends on several factors, including application start time, finish time, and required SLAs. In addition, workload models, whether an application is a long-running (24 × 7) web application, or a scientific workflow model of which its tasks need to be aware of precedence constraints, or applications based on IoT paradigm that is predominantly event-driven. Although one can optimise numerous scheduling parameters, many recent studies have focused on energy optimisation as a priority in Cloud data centre scheduling.

Chen et al. [37] propose energy-efficient scheduling in uncertain Cloud environments. They propose an interval number theory to define uncertainty, and a scheduling architecture manages this uncertainty in task scheduling. The proposed PRS1 scheduling algorithm based on proactive and reactive scheduling methods optimises energy in independent task scheduling. Similarly, Huang et al. [38] investigate energy-efficient scheduling for parallel workflow applications in

Cloud. Their EES algorithm tries to slack non-critical jobs to achieve power saving by exploiting the scheduling process's slack room. Energy-efficient scheduling using various heuristics for different application models has been a widely studied topic in literature [39]–[41].

Application of ML: A vast number of study explored application of ML in data centre scheduling. Some solutions rely on predictive models and then use them in scheduling algorithms, while other techniques model scheduling as a complete learning-based problem using Reinforcement learning (RL). Berrai et al. [42], adopt ML-based regression techniques to predict CPU load, power, SLAs and then use these in scheduling for better decisions. These solutions still use some level of heuristics with integrated prediction models. However, RL-based scheduling is designed to learn and take action in a data centre environment without explicit heuristics. Cheng et al. [43] proposed DRL-based provisioning and scheduling for application tasks in the data centre.

3) *Energy-aware Consolidation*: Cloud data centres are designed to handle the peak load to avoid potential SLA violations or overload conditions. Hence, the resources are oversubscribed to manage such an adverse situation. However, this oversubscription leads to resource underutilisation in general. It is estimated that Cloud data centres' utilisation level is around 50% on average. Underutilisation of resources is the main factor in the data centre's energy inefficiency as idle or lower utilised servers consume significant energy (up to 70% [44]). Thus, it is necessary to manage workloads under such oversubscribed and underutilised environments. To that end, consolidation has been a widely used technique to increase energy efficiency. It aims to bring the workloads (VMs and containers) from underutilised servers and consolidate them on fewer servers, thus allowing the remaining servers to be kept in sleep/shut down mode to save energy. Many challenges exist in consolidation, including maintaining VM affinity, avoiding overutilisation, minimising SLA violation, and reducing application downtime due to workload migrations.

Beloglazov et al. [35] proposed various heuristics to consolidate the workload and answer the question, including which VMs to migrate, where to migrate and when to migrate to reduce potential SLA Violation. Many other solutions have broadly focused on energy efficiency along with optimising different parameters (cost reduction, failure management, etc) while consolidating workloads in the data centre [45], [46].

Application of ML: ML-based solutions are predominantly used in consolidation [47], [48]. Hsieh et al. [48] studied VM consolidation to reduce power cost and increase QoS. They predict the utilisation of resources using the Gray-Markov-based model and use the information for consolidation. Similarly, the authors [47] also use prediction for consolidation. They predict memory and network usage and perform consolidation of VMs in a data centre along with CPU. Few researchers have also used RL in energy-aware consolidation [49], [50]. Basu et al. [50] proposed Megh—a system that learns to migrate VMs in the data centre using RL. It proposes the dimensionality reduction technique using

dimensional polynomial space with a sparse basis to minimise the state space in their problem. Their system has shown that it achieves better energy efficiency and cost reduction compared to existing heuristics.

4) *Power Capping*: Data centres are designed to handle peak power consumption based on the workload and cooling system requirements. Hence, in general, data centres are under-provisioned with power. This power capping on data centre servers restricts the amount of energy available to individual servers even though they can consume their maximum limit, thus providing the required speed for workloads [51]. Managing resources and workload effectively in these power-constrained environments is necessary. It is essential to avoid power inefficiencies in limited power allocated across servers to achieve power proportional computing [52].

In this regard, different power capping mechanisms at the Cloud data centre level are studied. The authors [53] proposed a fast decentralised power capping (DPC) technique to reduce latency and manage power at the individual server. Dynamo [54] is the power management system used by Facebook data centres, which has hierarchical power distribution. The lowest level leaf controller regulates power in a group of servers. This leaf controller, based on a high-bucket-first heuristic, determines the amount of energy to be reduced in each server to meet the power cap limits to which it is constrained. It also considers workload priorities and avoids potential performance degradation due to its power capping. Controlling peak power consumption is also a widely studied approach [55] by designing a feedback controller, which periodically reads system-level power and configures the highest power state of servers, keeping the server within its power budget. Authors in [56] studied optimal power allocation in servers, which accounts for several factors, including power-to-frequency, the arrival rate of jobs, and maximum and minimum server frequency configuration. They have shown that allocating full power may not always result in the highest speed as expected. Some techniques have also explored enabling data centre service providers to dynamically manage the power caps by participating in an open electricity market and achieve cost and energy efficiency [57].

Application of ML: Kumbhare et al. [58] propose a prediction based power over subscription in cloud data centres. they predictions of workload performance criticality and virtual machine (VM) resource utilisation and use this information to over subscribe the resources and increase overall utilisation. With Random Forest (RF) and Gradient Boosting (GB) models are used to predict the workload criticality and VM utilisation and use per VM power capping controller to limit its resource usage based on these predictions. However, due to the close interconnection between power capping effect on CPU speed, thermal dissipation and also the presence of heterogeneity in servers and workloads, data centre level power capping workload management is a difficult task to achieve [59] as compared to other energy efficiency methods that are discussed in this paper.

5) *Renewable Energy Management*: Data centres consume colossal energy and contribute significantly to greenhouse gas emissions (CO_2). Data centre service providers continuously increase renewable or green energy (solar, wind) usage with minimal carbon footprints to decarbonise the data centres. However, green energy usage in the data centre is extremely challenging due to its intermittent nature of availability. In contrast, the Cloud data centre needs an uninterrupted power supply since Cloud workloads tend to run 24×7 . Therefore, managing workloads under the uncertain availability of renewable energy is a challenging research problem.

Several resource management techniques explored maximising renewable energy in data centres. They include workload shifting and placement across geo-distributed data centres [60]–[62] based on their carbon efficiency. Besides, delaying job execution if an application can tolerate the QoS [63] and job dispatching or load balancing workloads to match the available renewable energy at different data centres [64] are some popular techniques in this regard.

Application of ML: ML-based algorithms are promising in renewable energy management, as predicting the available green energy based on an environmental condition is crucial in resource management tasks [65]–[67]. For instance, researchers from Google developed [68] carbon aware workload scheduling strategies for batch processing jobs by predicting the next day's available renewable energy from their energy sources. Similarly, Authors in [69] explores forecasting the carbon intensity of geographically distributed data centres and provides temporal shifting of workloads to minimize overall carbon footprints of workload execution. Along with prediction models, RL methods are also used to solve optimisation problems in increasing green energy usages in data centres by mapping renewable energy sources and physical machines [66].

IV. TAXONOMY OF THERMAL MANAGEMENT IN CDCS

Similar to energy management, thermal management techniques span from an individual server to data centres. A taxonomy on thermal management solutions is presented in Fig. 5. This section categorises these techniques into two broad categories, i.e., micro-level or single server level and macro-level or data centre-level thermal management techniques. We describe and review existing approaches and ML-based approaches used in these two categories.

A. Server Level

To achieve optimal performance, especially in modern chips with very high power densities, thermal constraints are the most critical challenges. Hence, it is essential to operate processors within the predefined Thermal Design Power (TDP) limit [70]. The servers consume an enormous amount of energy and dissipate it as heat. It is crucial to keep processor or CPU temperature within the TDP limit to avoid damage to the processor's silicon components, and prevent from catastrophic device failures. Modern rack servers reach peak temperatures up to 90-100°C.

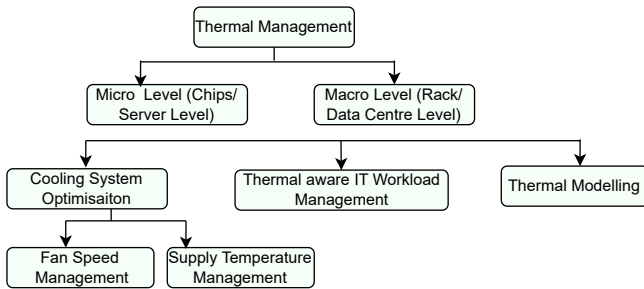


Fig. 5: Taxonomy of Thermal Management in Cloud Data Centres

Like DVFS in energy management, its corresponding thermal dissipation is regulated in servers by controlling the amount of power consumed. Dynamic Thermal Management (DTM) [71] is a popular thermal management technique at the individual server level which regulates Multiprocessors Systems-on-chip (MPSoCs) power consumption, and performance. This is done at the operating system level by closely communicating with underlying hardware interfaces. If a server's temperature is potentially exceeding the predefined TDP, the operating system takes actions by employing thermal throttling mechanisms to reduce energy consumption, thus reducing the CPU speed. Moreover, techniques like dynamic application scheduling [72], [73], onboard fan speed configuration [74] can also be employed for energy and thermal efficiency at the server level.

Application of ML: Recently, ML-based solutions have been applied to optimise temperature management at the individual server level. For example, Iranfar et al. [75] investigated how to proactively estimate the required number of active cores, operating frequency, and fan speed. Accordingly, the system is configured to achieve reduced power consumption and thus regulating corresponding server temperature. Although power consumption and CPU temperature are highly correlated, many other factors affect the thermal behaviour of servers including OS-level scheduling policies and compute-heavy applications. Therefore, analysing such resource management policies and workload behaviour through profiling, bench-marking, and then modelling through ML [32], is crucial for the design of future operating systems. As our focus is entirely on cloud data centres, we delve into more details on data centre level solutions in this regard.

B. Data Centre Level

A typical large-scale CDCs hosts thousands of servers. CDCs servers are arranged in rack-layout, where each rack (e.g., standard 42U rack) can accommodate 10-40 blade servers based on vendor-specific dimensions. This high density of equipment makes the data centre one of the highest-energy-density physical infrastructures. Dissipated heat from these rack servers can result in the data centre's ambient temperature reaching extremely high. Thus, cooling systems in data centres make sure that the data centre temperature

is within the threshold. Many approaches exist, optimising different parameters to reduce cooling energy. In this section, we review and describe data centre-level thermal management techniques.

1) *Cooling System Configurations:* Traditional rack layout data centres have a Computer Room Air Conditioning (CRAC) cooling system that blows cold air to the racks across the data centre (more details of cooling technologies can be found in Section IV-C). The entire cooling system efficiency requires multiple parameters to be configured in the design and operational phase. In the design phase, efficiency can be increased by better physical layout and vent designs to reduce heat re-circulations. While runtime cooling energy efficiency can be increased by fine-tuning the fan speeds of CRAC systems and cold air supply temperature, which determines the cooling system energy consumption [76]–[78]. In this section, we focus on runtime cooling system optimisation.

Fan Speed Management: Within the CRAC system, fans are used to regulate the airflow rate within the data centre. It is important to note that these fan speeds are separate from the onboard server's fan equipped to eject heat from CPU to the outside of the server cabinet. Increasing airflow requires higher fan speeds, thus consuming more energy. Hence, regulating fan speed optimally can save a significant amount of cooling power. However, this depends on the status of the data centre and its temperature level. Many researchers have proposed solutions to optimally configure the CRAC's fan speed based on cooling load [76], [79] by monitoring thermal load in the data centre and accordingly varying fan speeds dynamically to reduce energy consumption.

Supply Temperature Management: CRAC system blows cold air to racks through vented floor tiles in the data centre to take out dissipated heat. Passing colder air requires higher energy consumption as chillers in CRAC consume energy to supply cold air. Hence, the inaccurate configuration of supply air temperature significantly affects cooling energy costs in the data centre. For a safer operation, the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [80], recommends supply air temperature in the data centre to be in the range of 17-27 °C. Thus, it is beneficial to set the supply temperature closer to 27 °C. However, most data centres are overcooled as the supply temperature in the data centre is set to a much lower temperature conservatively, leaving energy inefficiencies in the cooling system. Setting a higher supply air temperature requires careful handling of peak temperature in data centres.

Many solutions have been proposed to raise the supply air temperature. Zhou et al. [81] have shown that significant power saving can be achieved when the workload is managed efficiently and allows the supply air temperature to be increased. In essence, to raise the supply air temperature, the data centre's peak temperature should be minimised. It can be done through various means, including thermal aware workload scheduling and avoiding thermal imbalance in the data centre.

Application of ML: In one of the earlier studies, Google used ML for cooling system optimisation in their data centres

[21]. The study employed a neural network framework to model and predict the Power Usage Effectiveness (PUE), which is the ratio of the total building energy usage to the IT energy usage. It used historical data from servers and other cooling systems, such as server IT load, server temperature, cooling set points, outside temperature, and more. The study analysed the effect of PUE on various configurations and provided feedback for system administrators to fine-tune the configurations efficiently, such as the number of dry coolers running, water pump speed, and the number of process water pumps. With the abilities of ML models to capture the complex nonlinear behaviour between parameters of different subsystems, they have high potential in modelling of the CDC cooling systems and help administrators increase the energy efficiency by adjusting the knobs.

2) *Thermal-aware IT Workload Management*: Thermal aware workload management include many sub tasks such as workload scheduling, workload consolidation, and workload dispersion, among others. These tasks significantly affects the thermal behaviours of a data centre. For instance, if the workload scheduling strategy results in peak temperature in the data centre, it generates a higher thermal load, thus increasing cooling costs. To address this, many researchers have proposed thermal-aware scheduling methods in Cloud data centres. Some solutions are proactive, which intends to avoid adverse temperature effects beforehand. In contrast, some scheduling policies follow reactive approaches. If a temperature violation is found, workloads are rescheduled to other nodes; however, the reactive scheduling method may result in higher QoS violations for applications due to rescheduling and migration. Mhedheb et al. [82] investigated load and thermal aware scheduling in Cloud that optimises temperature and load while scheduling tasks in data centres. Sun et al. [83] proposed thermal-aware scheduling of HPC jobs. They have used analytical models to estimate server temperature and model heat recirculation in the data centre. Proposed thermal-aware job assignment heuristics have shown increased performance with thermal balancing. Furthermore, authors in [84] have further extended thermal aware batch job scheduling across geo-distributed data centres.

Similarly, thermal-agnostic workload consolidation and dispersion triggers adverse temperature effects. Hence, balancing the workloads efficiently is necessary to achieve better efficiency. Consolidation is a widely used technique to optimise a computing system's energy consumption. However, aggressive consolidation leads to the creation of hotspots that further increases cooling cost. Hence, thermal-aware consolidation is necessary to balance the computing and cooling system energy consumption. A few studies have proposed solutions for this [85]–[87] to balance the temperature response due to workload placement during the workload consolidation. In contrast to consolidation, the workload dispersion technique aims to spread out workloads evenly across the data centre's servers [88], preventing peak utilisation in normal conditions. Although it minimises peak temperature, it significantly increases the computing system energy due to resource under-utilisation.

Hence, there should be a balance between consolidation and workload dispersion techniques to achieve cooling system efficiency.

Application of ML: Many of the existing works have employed machine-learning-based techniques in thermal-aware scheduling. Xiao et al. [89] presented a power and thermal-aware VM management framework based on machine learning, which relied on used Q-learning model to find optimal host configuration (power states) based on workload characteristics and cooling system's working state. The framework also enforced an efficient load-balancing policy to achieve a better trade-off between energy efficiency and performance. Similarly, many works have explored efficient distribution of application workload and also consolidation of VMs to increase resource utilisation and avoid thermal hotspots [90]–[92]. These works either develop temperature prediction model, aiding scheduling algorithms or develop controllers based on RL framework.

3) *Thermal Modelling*: Thermal modelling in data centres plays a vital role in resource management. Thermal modelling includes capturing thermal behaviour in a data centre and accurately estimating server temperature. Thermal models that predict accurately and fastly are useful aids in scheduling, configuring cooling systems and other resource management techniques. However, temperature prediction is a difficult problem. Server ambient temperature in a data centre depends on multiple factors, including CPU heat dissipation, inlet temperature and complex heat recirculation effects. There are mainly three types of thermal modelling techniques in data centres: (1) Computational Fluid Dynamic (CFD)-based models; (2) Analytical models; and (3) Predictive models. The CFD models accurately capture the room layouts, and heat recirculation effects and accurately estimates temperature in the data centre [93]–[95]. However, they are computationally expensive, and even a single calibration requires models to be run for multiple days. Hence, they are incapable of using them for fast online resource management decisions. On the other hand, analytical models depends on modelling data centre and workloads based on mathematical frameworks [83], [96]. They represent cooling, computing and workload elements with formal mathematical models and build a framework to establish relationships between all elements [83]. Although they are fast in temperature estimation, their accuracy is compromised due to their rigid static models.

Application of ML: ML-based predictive models use actual measurement data from the data centre to predict the accurate temperature of the server. These data-driven models, once trained, are accurate and quickly deliver the results in runtime. Moreover, they can automatically model the physical layout, air conditioning and the heat generated by Cloud data centres. Unlike CFD, where each of these needs to be modelled explicitly, this is a huge benefit. To that end, Wang et al. [97] proposed a server temperature prediction model using the Artificial Neural Network (ANN) based ML technique. Results have shown that it can accurately predict the temperature in data centres. In addition, some studies have explored using

machine learning models to identify temperature distribution [98] and to predict server inlet temperature [99].

The drawback of the ML-based model is that the model is only applicable to the data centre which the data is collected from. This means data need to be collected for each data centre extensively. However, this is not a massive disadvantage as such data need to be collected to monitor the data centres' health.

C. Cooling Technologies for Thermal Management in Data Centres

When servers/IT equipment uses electricity for their operations, the electrical energy is transferred as heat. This heat will be drawn across the server cabinet by the rear-mounted server fans allowing heat to transfer from the server's components to the outside ambient environment. Many technologies are employed to take out this heat from the data centre environment and keep the data centre's operational temperature within its threshold. These cooling technologies can be broadly categorised into two categories, including air and liquid cooling technologies.

1) *Air Cooling*: Air cooling is a widely used data centre cooling technology due to its inexpensive and flexible design and operational conveniences. In rack-layout-based data centres, the dissipated heat from servers is extracted from the cooling system's environment. The **Computer Room Air Conditioning (CRAC)** is a cooling system responsible for monitoring and managing the temperature in the data centre [100]. The CRAC blows cold air through the perforated tiles under the racks of a data centre. The cold air passes from the bottom to the top of the rack taking out the dissipated heat from rack equipment and this hot exhaust air is pushed to the intake of the CRAC units to the ceiling of the room, where it is taken out of the room. This allows the separation of hot exhaust air from cold inlet air. The CRAC unit then transfers the hot exhaust air via a coil to a fluid using refrigerant.

Many data centres also equip **Computer Room Air Handler (CRAH)**, where chilled water is used as fluid [100]. These fluids remove the heat from the data centre environment. The CRAC/CRAH continuously blow cold air using constant-speed fans, and this returns cold-air temperature, also called inlet temperature. It is configured to manage the dynamic thermal threshold in the data centre. It directly controls the cost of cooling in general. Lower the inlet temperature higher will be the cooling energy cost due to the increased energy required to transfer the lower temperature air from CRAC/CRAH. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [80], a leading technical Committee in cooling system technology, recommends that the device inlet be between 18-27°C for the safe operation of the environment. The design goal of any data centre operators will be to provide an inlet temperature close to 27 °C to reduce the cooling cost. However, the safer operation threshold should be maintained while configuring this parameter. Many works have looked into optimising this parameter using different techniques by minimising the peak temperature [96] by

balancing the workloads [92] and optimally configuring other parameters [101] of the cooling system.

Some modern systems also use **evaporative** [102] and **air side economisers/ free cooling** techniques [103]. In the evaporative technique, instead of fluid refrigerant, the hot air carried from the data centre is directly exposed to water. Water evaporates, taking out the heat from the hot air. Cooling towers are employed to dissipate the excess heat to the outside atmosphere. However, it doesn't require expensive CRAC or CRAH units but needs a large amount of water, a limiting factor in many data centre locations. On the other hand, air-side economisers or free cooling methods use outside free air for direct cooling instead of depending on the fluids to cool down the hot air extracted from CRAC/CRAH. This saves a huge amount of cooling costs. Nonetheless, these techniques vastly depend on the weather and geographical condition where the data centre is located, and thus they are used in limited computing infrastructures in practice.

2) *Liquid Cooling*: The recent advancement in data centre cooling technology has seen the adoption of liquid cooling as it is more efficient than air cooling, in general, [104]. The liquid cooling system also effectively avoids heat mix-up and heat re-circulation issues, which is a common problem in air cooling techniques.

Direct liquid cooling. In this system, liquid pipes are used to deliver liquid coolant directly to the heat sink present in the server's motherboards. The dissipated heat from the server is extracted to heat the chiller plant from these pipes, where the chilled water loop takes out the heat extracted from servers.

Immersion cooling. The computing system (servers and networking equipment is directly immersed in a non-conductive liquid. The liquid absorbs the heat and transfers it away from the components [105]. In some cases, equipment is arranged in isolated cabinets and immersed in tanks or cabinets are directly immersed in natural water habitats such as lakes/oceans. For instance, Microsoft has tested an underwater data centre with their project Natick [106], which allows them to operate the data centre in an energy-efficient manner by leveraging heat-exchange techniques with outside water. This technique is commonly used in submarines. This experimental project shows that immersion cooling is viable in large-scale computing systems with a group of servers sealed into large submarine cabinets.

Some other techniques have also been explored but are rarely used in large-scale settings, such as Dielectric fluid, where server components are coated with a non-conductive liquid. The heat is removed from the system by circulating liquid into direct contact with hot components, then through cool heat exchangers. Such methods are not widely adopted yet in practice. The common issue with rack-level liquid cooling is a lack of standardisation and specifications among multi-vendors. However, due to its energy efficiency compared to air cooling, it is expected that liquid cooling will become mainstream in future data centre cooling systems.




Fig. 6: Taxonomy of Integrated Energy and Thermal Management in Cloud Data Centres

V. TAXONOMY OF INTEGRATED ENERGY AND THERMAL MANAGEMENT IN CDCS

Traditionally cooling systems and computing systems are optimised individually. However, these two subsystems in the data centre are closely interdependent and optimising one system often have a opposite effect on others. Hence, the joint optimisation of two subsystems is beneficial, but it is challenging task that requires capturing complex dynamics of data centre workloads and physical environments. Fig. 6 shows a taxonomy of existing resource management solutions in integrated management of both computing and cooling energy.

Workload Management. A few studies have proposed solutions, including workload scheduling and cooling system optimisation as a multiobjective optimisation problem and accordingly configure different parameters to minimise energy consumption holistically [107], [108]. Other techniques include CRAC fan speed management by interplaying with IT load and its heat dissipation, configuring supply air temperature, and distributing the workload to minimise peak temperature, among many others.

Knobs Configurations. Wan et al. [109] studied holistic energy minimisation in data centres through a cross-layer optimisation framework for cooling and computing systems. This energy minimisation problem is formulated as a mixed-integer nonlinear programming problem. To solve this problem, the authors proposed a heuristic algorithm called JOINT, that dynamically configures parameters (such as server frequency, fan speed, and CRAC supply air temperature) based on workload demand and minimises computing and cooling system energy holistically.

Li et al. proposed [110] joint optimisation of computing and cooling systems for energy minimisation in data centres by modelling IT systems interactions (load distributions) and their corresponding thermal behaviour, i.e., heat transfer. The proposed analytical models for load distribution across rack servers minimise computing and cooling system energy, thereby configuring different knobs of two systems while ensuring the required throughput and resource constraints of workloads.

Power Budget Shifting. Power budget shifting is another important resource management technique in the Joint optimisation of these two systems. Using available power to trade between two systems in runtime can increase energy efficiency and resource utilisation. PowerTrade [111] is a technique that trades off data centre computing systems' idle power and cooling power with each other to reduce total power. Over-

provisioning is necessary for such conditions to accommodate extra workload and use excessive power obtained.

Application of ML: ML techniques have also been explored in the joint optimisation of computing and cooling systems. Recent advancements in RL have made it possible to learn different policies by interacting with the environments and learning from experience. RL techniques can be more adaptive and automatically understand the interdependence's of parameters. Ran et al. [112] used DRL and designed a hybrid action space that optimises the IT system and the airflow rate of the cooling system. Furthermore, the proposed control mechanism coordinates both the IT system's workload and cooling systems for energy efficiency. Similar techniques can be found in other studies [43], [113]. Careful design of state management, action, and rewards are important for applying RL techniques to data centres' holistic energy management.

VI. FUTURE RESEARCH DIRECTIONS

The sustainability in CDCs can be achieved by tackling some key issues that demand careful investigation and solutions. We need to fundamentally rethink how the data centres are currently managed, from hardware level optimisation to geo distributed data centre management. According to [11], if energy-aware approaches are implemented in CDCs, we can reduce total energy consumption in data centres up to 80% from the expected worst case scenario (Fig. 2). In the following, we identify key future research directions that should be pursued in this direction in order to reduce the energy footprints of CDCs and briefly explain them.

A. Standardisation and Tools for AI-centric RMS

One of the main barriers in adopting AI or ML solutions in data centre RMS is the lack of standardisation and tools. ML solutions need a lot of data. Currently, distributed systems, including Cloud systems, produce huge amounts of data from different computing layers. Standard methods and semantics are needed to collect, monitor, and interpret these data to accelerate the adoption of AI-centric models. Moreover, software tools and libraries need to be developed specifically for resource management systems, which will easily integrate policies into existing systems.

B. Hardware Software Co-design for ML-driven Resource Management

Computing servers and their components are tightly bound to operating systems that use simple rules to manage resources. This makes it hard to integrate new resource management policies that use ML to optimise hardware performance, because different vendors do not have common interfaces that can communicate with software. To solve this problem, we need a hardware-software co-design approach that allows us to develop and implement new resource management policies on hardware resources in an interoperable way.

C. Moving from "time-to-solution" to "Kw-to-solution"

Software development paradigms, platforms, and algorithms aim to enhance the execution speed of applications, but ignore their energy consumption. Hence, a paradigm shift is required from "time-to-solution" to "Kw-to-Solution" in software development and deployment. We also need new tools and programming constructs that enable software developers to assess and minimize the energy cost of their application logic, while preserving high speed. ML methods can offer valuable techniques for achieving this goal, such as learning energy-efficient code patterns, optimizing code performance, and adapting to different hardware configurations.

D. Resource Management in Emerging Cloud Workload Models using ML

Cloud computing is evolving from partially managed to fully managed services with application execution models like Serverless computing. Serverless computing lets us build applications with multiple stateless microservices. Cloud service providers handle the lifecycle of these microservices or stateless functions with guaranteed automatic scalability. This creates new challenges in pricing and managing thousands of stateless application services. ML-driven solutions, such as predicting user requests to cache "Hot" functions and reduce serverless function latency, or identifying resource interference among different user functions with classification methods, are some of the promising ways to address these challenges.

E. ML-Driven Holistic Resource Management

Cloud data centres consist of computing, networking, storage and cooling systems that are interdependent and crucial for ensuring service reliability. ML-driven resource management can detect these interdependencies and optimize the resources in a holistic manner to minimize energy consumption. A promising approach is to develop new algorithms and platforms that adjust parameters across different subsystems and balance tradeoffs.

F. Decarbonising Cloud Computing using ML

Cloud data centres are a major source of CO₂ emissions due to their dependence on fossil fuel-based energy sources. To decarbonise Cloud systems, many service providers are investing in renewable energy. However, the adoption of renewable energy sources is limited by their intermittent availability. Therefore, new solutions are needed to address the challenges of energy storage and workload management under uncertain energy supply. One promising direction is to use ML models to forecast the amount of renewable energy available at different Cloud data centre locations for a given time period. This prediction can enable the planning and execution of workloads in data centres that have more renewable energy, and thus reduce the reliance on fossil fuels.

G. Data-Driven Methods for Sustainable Multi-tier Computing Platforms

Multi-tier computing paradigms, such as Edge/Fog computing, have emerged to support IoT applications with distributed computations from the network edge to remote clouds. These paradigms pose new challenges for resource and application management, as they require low latency response and entail moving Cloud services from centralised locations to the network edge. Moreover, these paradigms involve more heterogeneous and energy-constrained environments than remote Clouds. Therefore, new solutions and approaches are needed for effective application and resource management under these conditions. ML methods can offer promising techniques for addressing these challenges, such as learning optimal resource allocation strategies, predicting workload patterns, and adapting to dynamic environments.

VII. CONCLUSIONS

Cloud computing platforms enable the development of highly connected resource-intensive applications, but they also require massive, heterogeneous, and complex data centres as their backbone infrastructure. Managing the energy and thermal aspects of such data centres is a challenging task, as the existing rule-based or heuristics solutions are not adequate to cope with the scale, heterogeneity, and dynamicity of the Cloud environment. Therefore, we need data-driven AI solutions that can leverage the data, learn from the environment, and make optimal resource management decisions. In this paper, we have explored leveraging AI-centric solutions for energy and thermal management in Cloud data centres. We have proposed a taxonomy for classifying different resource management techniques. We have also surveyed the state-of-the-art techniques and highlighted their strengths and limitations. Finally, we have suggested some promising future research directions

REFERENCES

- [1] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation computer systems*, vol. 25, no. 6, pp. 599–616, 2009.
- [2] A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, *et al.*, "Above the clouds: A Berkeley view of cloud computing," *University of California, Berkeley, Rep. UCB/EECS*, vol. 28, no. 13, p. 2009, 2009.
- [3] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer networks and ISDN systems*, vol. 30, no. 1-7, pp. 107–117, 1998.
- [4] Amazon, "Amazon Web Services."
- [5] Gartner, "Gartner forecasts worldwide public cloud revenue to grow 17 percent in 2020." <https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020>, 2019. [Online; accessed 10-Jan-2021].
- [6] Gartner, "Gartner forecasts worldwide public cloud end-user spending to grow 18 percent in 2021." <https://www.gartner.com/en/newsroom/press-releases/2020-11-17-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-grow-18-percent-in-2021>, 2020. [Online; accessed 10-Jan-2021].

- [7] R. Buyya, S. Ilager, and P. Arroba, "Energy-efficiency and sustainability in new generation cloud computing: A vision and directions for integrated management of data centre resources and workloads," *Software: Practice and Experience*, vol. 54, no. 1, pp. 24–38, 2024.
- [8] S. Maybury, "How much Energy does your Data Centre Use?.." https://www.metronode.com.au/energy_usage/, 2017. [Online; accessed 05-Jan-2021].
- [9] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United states data center energy usage report," 2016.
- [10] J. Koomey, "Growth in data center electricity use 2005 to 2010," *A report by Analytical Press, completed at the request of The New York Times*, vol. 9, 2011.
- [11] A. S. Andrae and T. Edler, "On global electricity usage of communication technology: trends to 2030," *Challenges*, vol. 6, no. 1, pp. 117–157, 2015.
- [12] P. Johnson and T. Marker, "Data centre energy efficiency product profile," *Pitt & Sherry, report to equipment energy efficiency committee (E3) of The Australian Government Department of the Environment, Water, Heritage and the Arts (DEWHA)*, 2009.
- [13] H. Viswanathan, E. K. Lee, and D. Pompili, "Self-organizing sensing infrastructure for autonomic management of green datacenters," *IEEE Network*, vol. 25, no. 4, pp. 34–40, 2011.
- [14] D. Minoli, K. Sohrawy, and B. Occhiogrosso, "Iot considerations, requirements, and architectures for smart buildings—energy optimization and next-generation building management systems," *IEEE Internet of Things Journal*, vol. 4, no. 1, pp. 269–283, 2017.
- [15] Q. Liu, Y. Ma, M. Alhussein, Y. Zhang, and L. Peng, "Green data center with iot sensing and cloud-assisted smart temperature control system," *Computer Networks*, vol. 101, pp. 104–112, 2016.
- [16] S. Saha and A. Majumdar, "Data centre temperature monitoring with esp8266 based wireless sensor network and cloud based dashboard with real time alert system," in *2017 Devices for Integrated Circuit (DevIC)*, pp. 307–310, IEEE, 2017.
- [17] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green ai," *arXiv preprint arXiv:1907.10597*, 2019.
- [18] D. Amodei and D. Hernandez, "Ai and compute," 2018. <https://blog.openai.com/ai-and-compute>.
- [19] D. Jeff, "ML for system, system for ML, keynote talk in Workshop on ML for Systems, NIPS," 2018.
- [20] R. Bianchini, M. Fontoura, E. Cortez, A. Bonde, A. Muzio, A.-M. Constantin, T. Moscibroda, G. Magalhaes, G. Bablani, and M. Russinovich, "Toward ml-centric cloud platforms," *Communications of the ACM*, vol. 63, no. 2, pp. 50–59, 2020.
- [21] J. Gao, "Machine learning applications for data center optimization," *Google White Paper*, 2014.
- [22] W. Torell, K. Brown, and V. Avelar, "The unexpected impact of raising data center temperatures," *Write paper 221, Revision*, 2015.
- [23] V. Venkatachalam and M. Franz, "Power reduction techniques for microprocessor systems," *ACM Computing Surveys (CSUR)*, vol. 37, no. 3, pp. 195–237, 2005.
- [24] E.-Y. Chung, L. Benini, and G. De Micheli, "Dynamic power management using adaptive learning tree," in *Proceedings of the 1999 IEEE/ACM International Conference on Computer-Aided Design. Digest of Technical Papers (Cat. No. 99CH37051)*, pp. 274–279, IEEE, 1999.
- [25] Y. Lu, D. Wu, B. He, X. Tang, J. Xu, and M. Guo, "Rank-aware dynamic migrations and adaptive demotions for dram power management," *IEEE Transactions on Computers*, vol. 65, no. 1, pp. 187–202, 2015.
- [26] R. A. Bridges, N. Imam, and T. M. Mintz, "Understanding GPU power: A survey of profiling, modeling, and simulation methods," *ACM Computing Surveys*, vol. 49, no. 3, 2016.
- [27] G. Von Laszewski, L. Wang, A. J. Younge, and X. He, "Power-aware scheduling of virtual machines in dvfs-enabled clusters," in *Proceedings of the 2009 IEEE International Conference on Cluster Computing and Workshops*, pp. 1–10, IEEE, 2009.
- [28] P. Arroba, J. M. Moya, J. L. Ayala, and R. Buyya, "Dynamic voltage and frequency scaling-aware dynamic consolidation of virtual machines for energy efficient cloud data centers," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 10, p. e4067, 2017.
- [29] M. Safari and R. Khorsand, "Energy-aware scheduling algorithm for time-constrained workflow tasks in dvfs-enabled cloud environment," *Simulation Modelling Practice and Theory*, vol. 87, pp. 311–326, 2018.
- [30] S. Wang, Z. Qian, J. Yuan, and I. You, "A dvfs based energy-efficient tasks scheduling in a data center," *IEEE Access*, vol. 5, pp. 13090–13102, 2017.
- [31] J.-G. Park, N. Dutt, and S.-S. Lim, "Ml-gov: A machine learning enhanced integrated cpu-gpu dvfs governor for mobile gaming," in *Proceedings of the 15th IEEE/ACM Symposium on Embedded Systems for Real-Time Multimedia*, pp. 12–21, 2017.
- [32] S. Ilager, R. Wankar, R. Kune, and R. Buyya, "Gpu paas computation model in aneka cloud computing environments," *Smart Data: State-of-the-Art Perspectives in Computing and Applications*, p. 19, 2019.
- [33] X. Xu, W. Dou, X. Zhang, and J. Chen, "Enreal: An energy-aware resource allocation method for scientific workflow executions in cloud environment," *IEEE Transactions on Cloud Computing*, vol. 4, no. 2, pp. 166–179, 2015.
- [34] H. Li, J. Li, W. Yao, S. Nazarian, X. Lin, and Y. Wang, "Fast and energy-aware resource provisioning and task scheduling for cloud systems," in *Proceedings of the 18th International Symposium on Quality Electronic Design (ISQED)*, pp. 174–179, IEEE, 2017.
- [35] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," *Future Generation Computer Systems*, vol. 28, no. 5, pp. 755–768, 2012.
- [36] M. Dabbagh, B. Hamdaoui, M. Guizani, and A. Rayes, "Energy-efficient resource allocation and provisioning framework for cloud data centers," *IEEE Transactions on Network and Service Management*, vol. 12, no. 3, pp. 377–391, 2015.
- [37] H. Chen, X. Zhu, H. Guo, J. Zhu, X. Qin, and J. Wu, "Towards energy-efficient scheduling for real-time tasks under uncertain cloud computing environment," *Journal of Systems and Software*, vol. 99, pp. 20–35, 2015.
- [38] Q. Huang, S. Su, J. Li, P. Xu, K. Shuang, and X. Huang, "Enhanced energy-efficient scheduling for parallel applications in cloud," in *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012)*, pp. 781–786, IEEE, 2012.
- [39] Y. Ding, X. Qin, L. Liu, and T. Wang, "Energy efficient scheduling of virtual machines in cloud with deadline constraint," *Future Generation Computer Systems*, vol. 50, pp. 62–74, 2015.
- [40] R. N. Calheiros and R. Buyya, "Energy-efficient scheduling of urgent bag-of-tasks applications in clouds through dvfs," in *Proceedings of the 6th IEEE international conference on cloud computing technology and science*, pp. 342–349, IEEE, 2014.
- [41] C. Ghribi, M. Hadji, and D. Zeghlache, "Energy efficient vm scheduling for cloud data centers: Exact allocation and migration algorithms," in *Proceedings of the 13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*, pp. 671–678, IEEE, 2013.
- [42] J. L. Berral, Í. Goiri, R. Nou, F. Julià, J. Guitart, R. Gavaldà, and J. Torres, "Towards energy-aware scheduling in data centers using machine learning," in *Proceedings of the 1st International Conference on energy-Efficient Computing and Networking*, pp. 215–224, 2010.
- [43] M. Cheng, J. Li, and S. Nazarian, "Drl-cloud: Deep reinforcement learning-based resource provisioning and task scheduling for cloud service providers," in *Proceedings of the 23rd Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 129–134, IEEE, 2018.
- [44] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Zomaya, "A taxonomy and survey of energy-efficient data centers and cloud computing systems," in *Advances in Computers*, vol. 82, pp. 47–111, Elsevier, 2011.
- [45] S. F. Piraghaj, A. V. Dastjerdi, R. N. Calheiros, and R. Buyya, "A framework and algorithm for energy efficient container consolidation in cloud data centers," in *Proceedings of the 2015 IEEE International Conference on Data Science and Data Intensive Systems*, pp. 368–375, IEEE, 2015.
- [46] M. H. Ferdous, M. Murshed, R. N. Calheiros, and R. Buyya, "Virtual machine consolidation in cloud data centers using aco metaheuristic," in *Proceedings of the European conference on parallel processing*, pp. 306–317, Springer, 2014.
- [47] N. T. Hieu, M. Di Francesco, and A. Ylä-Jääski, "Virtual machine consolidation with multiple usage prediction for energy-efficient cloud data centers," *IEEE Transactions on Services Computing*, vol. 13, no. 1, pp. 186–199, 2017.
- [48] S.-Y. Hsieh, C.-S. Liu, R. Buyya, and A. Y. Zomaya, "Utilization-prediction-aware virtual machine consolidation approach for energy-efficient cloud data centers," *Journal of Parallel and Distributed Computing*, vol. 139, pp. 99–109, 2020.

- [49] F. Farahnakian, P. Liljeberg, and J. Plosila, "Energy-efficient virtual machines consolidation in cloud data centers using reinforcement learning," in *2014 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*, pp. 500–507, IEEE, 2014.
- [50] D. Basu, X. Wang, Y. Hong, H. Chen, and S. Bressan, "Learn-as-you-go with megh: Efficient live migration of virtual machines," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 8, pp. 1786–1801, 2019.
- [51] A. A. Bhattacharya, D. Culler, A. Kansal, S. Govindan, and S. Sankar, "The need for speed and stability in data center power capping," *Sustainable Computing: Informatics and Systems*, vol. 3, no. 3, pp. 183–193, 2013.
- [52] P. Petoumenos, L. Mukhanov, Z. Wang, H. Leather, and D. S. Nikolopoulos, "Power capping: What works, what does not," in *2015 IEEE 21st International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 525–534, IEEE, 2015.
- [53] R. Azimi, M. Badiei, X. Zhan, N. Li, and S. Reda, "Fast decentralized power capping for server clusters," in *HPCA*, pp. 181–192, 2017.
- [54] Q. Wu, Q. Deng, L. Ganesh, C.-H. Hsu, Y. Jin, S. Kumar, B. Li, J. Meza, and Y. J. Song, "Dynamo: Facebook's data center-wide power management system," *ACM SIGARCH Computer Architecture News*, vol. 44, no. 3, pp. 469–480, 2016.
- [55] C. Lefurgy, X. Wang, and M. Ware, "Power capping: a prelude to power shifting," *Cluster Computing*, vol. 11, no. 2, pp. 183–195, 2008.
- [56] A. Gandhi, M. Harchol-Balter, R. Das, and C. Lefurgy, "Optimal power allocation in server farms," *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 1, pp. 157–168, 2009.
- [57] H. Chen, C. Hankendi, M. C. Caramanis, and A. K. Coskun, "Dynamic server power capping for enabling data center participation in power markets," in *Proceedings of the 2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pp. 122–129, IEEE, 2013.
- [58] A. G. Kumbhare, R. Azimi, I. Manousakis, A. Bonde, F. Frujeri, N. Mahalingam, P. A. Misra, S. A. Javadi, B. Schroeder, M. Fontoura, and R. Bianchini, "Prediction-Based power oversubscription in cloud platforms," in *Proceedings of the 2021 USENIX Annual Technical Conference (USENIX ATC 21)*, (Berkeley, CA, USA), pp. 473–487, USENIX Association, July 2021.
- [59] H. Zhang and H. Hoffmann, "Maximizing performance under a power cap: A comparison of hardware, software, and hybrid techniques," *ACM SIGPLAN Notices*, vol. 51, no. 4, pp. 545–559, 2016.
- [60] M. Xu and R. Buyya, "Managing renewable energy and carbon footprint in multi-cloud computing environments," *Journal of Parallel and Distributed Computing*, vol. 135, pp. 191–202, 2020.
- [61] A. Khosravi, L. L. Andrew, and R. Buyya, "Dynamic vm placement method for minimizing energy and carbon cost in geographically distributed cloud data centers," *IEEE Transactions on Sustainable Computing*, vol. 2, no. 2, pp. 183–196, 2017.
- [62] U. Mandal, M. F. Habib, S. Zhang, B. Mukherjee, and M. Tornatore, "Greening the cloud using renewable-energy-aware service migration," *IEEE network*, vol. 27, no. 6, pp. 36–43, 2013.
- [63] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Greenhadoop: leveraging green energy in data-processing frameworks," in *Proceedings of the 7th ACM european conference on Computer Systems*, pp. 57–70, 2012.
- [64] Y. Zhang, Y. Wang, and X. Wang, "Greenware: Greening cloud-scale data centers to maximize the use of renewable energy," in *Proceedings of the ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*, pp. 143–164, Springer, 2011.
- [65] J.-P. Lai, Y.-M. Chang, C.-H. Chen, and P.-F. Pai, "A survey of machine learning models in renewable energy predictions," *Applied Sciences*, vol. 10, no. 17, p. 5975, 2020.
- [66] J. Gao, H. Wang, and H. Shen, "Smartly handling renewable energy instability in supporting a cloud datacenter," in *Proceedings of the 2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 769–778, IEEE, 2020.
- [67] L. Lin and A. A. Chien, "Adapting datacenter capacity for greener datacenters and grid," in *Proceedings of the 14th ACM International Conference on Future Energy Systems*, pp. 200–213, 2023.
- [68] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care, S. Talukdar, E. Mullen, K. Smith, M. Cottman, and W. Cirne, "Carbon-aware computing for datacenters," *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2023.
- [69] P. Wiesner, I. Behnke, D. Scheinert, K. Gontarska, and L. Thamsen, "Let's wait awhile: how temporal workload shifting can reduce carbon emissions in the cloud," in *Proceedings of the 22nd International Middleware Conference*, Middleware '21, (New York, NY, USA), p. 260–272, Association for Computing Machinery, 2021.
- [70] S. Pagani, P. S. Manoj, A. Jantsch, and J. Henkel, "Machine learning for power, energy, and thermal management on multicore processors: A survey," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 1, pp. 101–116, 2018.
- [71] D. Shin, S. W. Chung, E.-Y. Chung, and N. Chang, "Energy-optimal dynamic thermal management: Computation and cooling power co-optimization," *IEEE Transactions on Industrial Informatics*, vol. 6, no. 3, pp. 340–351, 2010.
- [72] R. Ayoub, K. Indukuri, and T. S. Rosing, "Temperature aware dynamic workload scheduling in multisocket cpu servers," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 30, no. 9, pp. 1359–1372, 2011.
- [73] H. F. Sheikh, I. Ahmad, Z. Wang, and S. Ranka, "An overview and classification of thermal-aware scheduling techniques for multi-core processing systems," *Sustainable Computing: Informatics and Systems*, vol. 2, no. 3, pp. 151–169, 2012.
- [74] Z. Wang, C. Bash, N. Tolia, M. Marwah, X. Zhu, and P. Ranganathan, "Optimal fan speed control for thermal management of servers," in *Proceedings of the International Electronic Packaging Technical Conference and Exhibition*, vol. 43604, pp. 709–719, 2009.
- [75] A. Iranfar, F. Terraneo, G. Csordas, M. Zapater, W. Fornaciari, and D. Atienza, "Dynamic thermal management with proactive fan speed control through reinforcement learning," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 418–423, IEEE, 2020.
- [76] W. Zhang, Y. Wen, Y. W. Wong, K. C. Toh, and C.-H. Chen, "Towards joint optimization over ict and cooling systems in data centre: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 1596–1616, 2016.
- [77] A. H. Khalaj and S. K. Halgamuge, "A review on efficient thermal management of air-and liquid-cooled data centers: From chip to the cooling system," *Applied energy*, vol. 205, pp. 1165–1188, 2017.
- [78] C. Nadjahi, H. Louahlia, and S. Lemasson, "A review of thermal management and innovative cooling strategies for data center," *Sustainable Computing: Informatics and Systems*, vol. 19, pp. 14–28, 2018.
- [79] M. Tian, *Energy Optimization by Fan Speed Control for Data Centers*. PhD thesis, The George Washington University, 2019.
- [80] ASHRAE, "American society of heating, refrigerating and air-conditioning engineers," 2018. URL: <http://tc0909.ashraetsc.org/>.
- [81] R. Zhou, Z. Wang, C. E. Bash, A. McReynolds, C. Hoover, R. Shih, N. Kumari, and R. K. Sharma, "A holistic and optimal approach for data center cooling management," in *Proceedings of the 2011 American Control Conference*, pp. 1346–1351, IEEE, 2011.
- [82] Y. Mhedheb, F. Jrad, J. Tao, J. Zhao, J. Kołodziej, and A. Streit, "Load and thermal-aware vm scheduling on the cloud," in *Proceedings of the International Conference on Algorithms and Architectures for Parallel Processing*, pp. 101–114, Springer, 2013.
- [83] H. Sun, P. Stolf, and J.-M. Pierson, "Spatio-temporal thermal-aware scheduling for homogeneous high-performance computing datacenters," *Future Generation Computer Systems*, vol. 71, pp. 157–170, 2017.
- [84] M. Polverini, A. Cianfrani, S. Ren, and A. V. Vasilakos, "Thermal-aware scheduling of batch jobs in geographically distributed data centers," *IEEE Transactions on cloud computing*, vol. 2, no. 1, pp. 71–84, 2013.
- [85] E. K. Lee, H. Viswanathan, and D. Pompili, "Vmap: Proactive thermal-aware virtual machine allocation in hpc cloud datacenters," in *2012 19th International Conference on High Performance Computing*, pp. 1–10, IEEE, 2012.
- [86] S. Ilager, K. Ramamohanarao, and R. Buyya, "Etas: Energy and thermal-aware dynamic virtual machine consolidation in cloud data center with proactive hotspot mitigation," *Concurrency and Computation: Practice and Experience*, vol. 0, no. 0, p. e5221, 2019.
- [87] J. V. Wang, C.-T. Cheng, and C. K. Tse, "A thermal-aware vm consolidation mechanism with outage avoidance," *Software: Practice and Experience*, vol. 49, no. 5, pp. 906–920, 2019.

- [88] H. Shamalizadeh, L. Almeida, S. Wan, P. Amaral, S. Fu, and S. Prabh, "Optimized thermal-aware workload distribution considering allocation constraints in data centers," in *2013 IEEE international conference on green computing and communications and IEEE internet of things and IEEE cyber, physical and social computing*, pp. 208–214, IEEE, 2013.
- [89] P. Xiao, Z. Ni, D. Liu, and Z. Hu, "A power and thermal-aware virtual machine management framework based on machine learning," *Cluster Computing*, vol. 24, pp. 2231–2248, 2021.
- [90] S. Akbar, R. Li, M. Waqas, and A. Jan, "Server temperature prediction using deep neural networks to assist thermal-aware scheduling," *Sustainable Computing: Informatics and Systems*, vol. 36, p. 100809, 2022.
- [91] S. S. Gill, S. Tuli, A. N. Toosi, F. Cuadrado, P. Garraghan, R. Bahsoon, H. Lutfiyya, R. Sakellariou, O. Rana, S. Dustdar, *et al.*, "Thermosim: Deep learning based framework for modeling and simulation of thermal-aware resource management for cloud computing environments," *Journal of Systems and Software*, vol. 166, p. 110596, 2020.
- [92] S. Ilager, K. Ramamohanarao, and R. Buyya, "Thermal prediction for efficient energy management of clouds using machine learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 5, pp. 1044–1056, 2020.
- [93] J. Choi, Y. Kim, A. Sivasubramaniam, J. Srebric, Q. Wang, and J. Lee, "A cfd-based tool for studying temperature in rack-mounted servers," *IEEE Transaction on Computers*, vol. 57, no. 8, pp. 1129–1142, 2008.
- [94] R. Romadhon, M. Ali, A. M. Mahdzir, and Y. A. Abakr, "Optimization of cooling systems in data centre by computational fluid dynamics model and simulation," in *2009 Innovative Technologies in Intelligent Systems and Industrial Applications*, pp. 322–327, IEEE, 2009.
- [95] A. Almoli, A. Thompson, N. Kapur, J. Summers, H. Thompson, and G. Hannah, "Computational fluid dynamic investigation of liquid rack cooling in data centres," *Applied energy*, vol. 89, no. 1, pp. 150–155, 2012.
- [96] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 11, pp. 1458–1472, 2008.
- [97] L. Wang, G. von Laszewski, F. Huang, J. Dayal, T. Frulani, and G. Fox, "Task scheduling with ann-based temperature prediction in a data center: a simulation-based study," *Engineering with Computers*, vol. 27, no. 4, pp. 381–391, 2011.
- [98] Y. Tarutani, K. Hashimoto, G. Hasegawa, Y. Nakamura, T. Tamura, K. Matsuda, and M. Matsuoka, "Temperature distribution prediction in data centers for decreasing power consumption by machine learning," in *Proceedings of the 7th IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pp. 635–642, IEEE, 2015.
- [99] R. Lloyd and M. Rebow, "Data driven prediction model (ddpm) for server inlet temperature prediction in raised-floor data centers," in *Proceedings of the 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, pp. 716–725, IEEE, 2018.
- [100] B. Fakhim, M. Behnia, S. Armfield, and N. Srinarayana, "Cooling solutions in an operational data centre: A case study," *Applied Thermal Engineering*, vol. 31, no. 14–15, pp. 2279–2291, 2011.
- [101] E. K. Lee, H. Viswanathan, and D. Pompili, "Proactive thermal-aware resource management in virtualized hpc cloud datacenters," *IEEE Transactions on Cloud Computing*, vol. 5, no. 2, pp. 234–248, 2015.
- [102] B. Porumb, P. Ungureşan, L. F. Tutunaru, A. Şerban, and M. Bălan, "A review of indirect evaporative cooling operating conditions and performances," *Energy Procedia*, vol. 85, pp. 452–460, 2016.
- [103] H. Zhang, S. Shao, H. Xu, H. Zou, and C. Tian, "Free cooling of data centers: A review," *Renewable and Sustainable Energy Reviews*, vol. 35, pp. 171–182, 2014.
- [104] T. Gao, S. Shao, Y. Cui, B. Espiritu, C. Ingalz, H. Tang, and A. Heydari, "A study of direct liquid cooling for high-density chips and accelerators," in *Proceedings of the 16th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, pp. 565–573, IEEE, 2017.
- [105] A. Capozzoli and G. Primiceri, "Cooling systems in data centers: state of art and emerging technologies," *Energy Procedia*, vol. 83, pp. 484–493, 2015.
- [106] B. Cutler, S. Fowers, J. Kramer, and E. Peterson, "Dunking the data center," *IEEE Spectrum*, vol. 54, no. 3, pp. 26–31, 2017.
- [107] W. Zhang, Y. Wen, Y. W. Wong, K. C. Toh, and C.-H. Chen, "Towards joint optimization over ict and cooling systems in data centre: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 1596–1616, 2016.
- [108] X. Li, P. Garraghan, X. JIANG, Z. Wu, and J. Xu, "Holistic Virtual Machine Scheduling in Cloud Datacenters towards Minimizing Total Energy," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 6, pp. 1–1, 2017.
- [109] J. Wan, X. Gui, R. Zhang, and L. Fu, "Joint cooling and server control in data centers: A cross-layer framework for holistic energy minimization," *IEEE Systems Journal*, vol. 12, no. 3, pp. 2461–2472, 2017.
- [110] S. Li, H. Le, N. Pham, J. Heo, and T. Abdelzaher, "Joint optimization of computing and cooling energy: Analytic model and a machine room case study," in *Proceedings of the 2012 IEEE 32nd International Conference on Distributed Computing Systems*, pp. 396–405, IEEE, 2012.
- [111] F. Ahmad and T. Vijaykumar, "Joint optimization of idle and cooling power in data centers while maintaining response time," *ACM Sigplan Notices*, vol. 45, no. 3, pp. 243–256, 2010.
- [112] Y. Ran, H. Hu, X. Zhou, and Y. Wen, "Deepee: Joint optimization of job scheduling and cooling control for data center energy efficiency using deep reinforcement learning," in *Proceedings of the 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pp. 645–655, IEEE, 2019.
- [113] A. Mirhoseini, H. Pham, Q. V. Le, B. Steiner, R. Larsen, Y. Zhou, N. Kumar, M. Norouzi, S. Bengio, and J. Dean, "Device placement optimization with reinforcement learning," in *Proceedings of the 34th International Conference on Machine Learning*, pp. 2430–2439, JMLR.org, 2017.

The Interplay of Learning Analytics and Artificial Intelligence

Jelena Jovanovic
0000-0002-1904-0446
University of Belgrade, Faculty of
Organisational Sciences, Jove Ilica
154, Belgrade, Serbia
Email:
jelena.jovanovic@fon.bg.ac.rs

Abstract—The widespread use of digital systems and tools in education has opened up opportunities for collecting, measuring, and analysing data about user (learner, teacher) interactions with a variety of learning resources and activities, with the ultimate objective of better understanding learning and advancing both learning outcomes and the overall learning experience. This promise motivated the development of Learning Analytics (LA) as a research and practical field and the use of insights derived from learning trace data for evidence-based decision making in a variety of educational settings. While LA has made a significant contribution to better understanding of learning and the environments in which it takes place, many open questions and challenges remain. Furthermore, new opportunities and challenges continue to emerge with the ever-changing modalities of teaching and learning, the latest of which are associated with the rapid development and accessibility of Artificial Intelligence (AI). Taking the cyclical model of LA as its exploration framework, this paper examines how key components of the LA model – namely data, methods, and actions – relate to and may benefit from the latest developments in AI, and especially Generative AI. Aiming for evidence-based analysis and discussion of the interplay between LA and AI, the paper relies on the latest empirical research in LA and the related research fields of AI in Education and Educational Data Mining.

Index Terms—Learning Analytics, Artificial Intelligence in Education, Generative AI.

I. INTRODUCTION

THE educational landscape is undergoing a continuous digital transformation. Online and blended learning modalities are flourishing, and a vast array of software tools and gadgets are now commonplace in classrooms. These advancements allow for the unobtrusive gathering of data about learners' interactions with learning resources and other participants in the educational process. This wealth of data provides a rich foundation for understanding learning and advancing learning outcomes and the overall educational experience. Different approaches have emerged to achieve these objectives, ranging from fully automated systems aimed at personalising learning according to individual learners' needs and preferences to those that provide learners with

information—such as analytics, recommendations, and pedagogical scaffolds—empowering them to take initiative and adapt their learning pathways on their own. This paper focuses on the latter group of approaches, which emphasise user agency and adaptable learning processes and are central to the field of Learning Analytics (LA).

The recent rapid advancements and adoption of Artificial Intelligence (AI) has opened new opportunities and challenges in educational settings. Generative AI, with its advanced capabilities, promises to significantly impact how educational content is created, delivered, and used. The field of LA, with its established methodologies for studying learning, is well-positioned to systematically explore and understand the benefits and drawbacks of incorporating (Generative) AI into education.

Set against this backdrop, this paper aims to achieve two objectives. First, it introduces LA, highlighting its iterative nature and the key elements of the LA process. Second, it explores the interplay between LA and AI, by focusing on how LA can enhance our understanding of AI in education and how the LA process and its key components may benefit from advancements in (Generative) AI. By examining these dynamics, the paper aims to demonstrate how AI, and especially Generative AI, may empower LA to keep pace with the rapidly changing educational realm and stay true to its mission of understanding and advancing learning. In doing so, the paper relies on published empirical research in LA and closely related fields of AI in Education and Educational Data Mining. This evidence-based approach, inherent to LA, distinguishes the current paper from recent publications that discuss the opportunities and challenges of (Generative) AI for LA, and education more broadly, from a more hypothetical perspective.

II. LEARNING ANALYTICS

A. What is Learning Analytics?

Learning analytics is defined as the “measurement, collection, analysis and reporting of data about learners and their

contexts, for purposes of understanding and optimising learning and the environments in which it occurs” [1]. For a better understanding of LA, it is necessary to unpack this rather compact definition and highlight the key distinguishing features of LA as a research and practical field.

First, *data* are at the centre of any LA effort. LA uses a wide variety of data types and sources such as log data, self-reports, messages exchanged in distinct kinds of online communication channels, sensory data, etc. Amidst this variety of data types and sources, *learning traces* - also referred to as trace data or learning logs - remain the primary type of data in LA. Learning traces are data about learners’ interactions with different (digital) learning resources, (online) learning activities, as well as other learners and teachers (e.g., communication in online discussion forums). The main advantage of learning traces compared to data traditionally used in educational research (e.g., surveys and think aloud protocols) is that learning traces can be collected seamlessly during the learning process, without putting any additional burden on learners and teachers. The continuous increase in the number and variety of software platforms and tools used in the learning process, as well as the continuously increasing adoption of online and blended learning both in formal and non-formal education, make learning trace data more and more available. This trend positively reflects on the relevance and the adoption of LA in practice.

In addition to learning logs collected in the context of online and blended learning, the collection of learning-related data in traditional classrooms and physical spaces in general, attracts more and more interest from LA researchers. This is made possible by the increasing availability of sensors (e.g., cameras, microphones, location-tracking sensors) that allow for measuring and collecting data about learners’ interactions with a variety of physical objects used in learning, as well as data about mutual interactions of learners and teachers in different situations of collaborative learning. The collection and combined use of data from multiple sources, as well as advanced analytics such data enable, are in the focus of a sub-field of LA known as Multimodal LA [2].

Another key construct in the definition of LA that requires further explanation is *optimization of learning and the environment in which learning takes place*, which is stated as one of the main objectives of the field. It is important to highlight that the term optimization in this context does not imply automatic adaptation of the learning process to a particular learner (e.g., automated personalization of learning), as is the case in closely related fields of Artificial Intelligence in Education and Educational Data Mining. In LA, optimization means that the results of analytics, such as insights about a learning process or recommendations, are communicated to learners and/or teachers, and it is left to them to decide how to act on the feedback received. Acting on the feedback in case of mature learners may take the form of making adjustments to one’s own learning approach, in accordance with the information and recommendations received. In the case of young learners, feedback is typically directed to the teacher to

help them choose pedagogical interventions to better support their students. Simply put, in LA, it is important to include humans (students, teachers, parents, etc.) in the process of adaptation and improvement of learning, the concept often referred to as human-at-the-centre. This is in accordance with one of the most prevalent learning approaches in LA, namely self-regulated learning (SRL), which postulates that learner is an active agent who, in a learning process, first defines their goals, then chooses learning strategies and tactics to achieve those goals, and while acting in the direction of the goals, continuously monitors and evaluates their progress and adjusts the chosen strategies and tactics accordingly [3]. The primary role of LA is to support the learner at all stages of the learning process, providing evidence-based insights, recommendations, and guidelines. Furthermore, such an approach gives teachers the sense of being in control of their teaching work (instead of being replaced through automation), which facilitates technology adoption.

Finally, it is necessary to clarify the meaning of *learning context* in the LA definition, considering that this term has been assigned a variety of meanings in educational research and practice. In LA, learning context is often described as a specific combination of internal and external factors that may affect learning [4]. Here, a learner is considered the reference point, meaning that internal factors include everything that constitutes the internal state of the learner, such as emotional state, motivation, prior knowledge, cognitive load, etc. On the other hand, external factors include all that may affect learning and is external to the learner, i.e., the learner does not have direct control over (e.g., pedagogical design of the course, specific pedagogical approach of the teacher, class schedule, etc.).

All the above suggests that LA is an interdisciplinary field, at the intersection of fields focused on learning (pedagogy, educational psychology, educational technologies), analytics (computer science, statistics, artificial intelligence), and human-centred design (human-computer interaction).

B. Learning Analytics Cycle

Learning Analytics can be viewed as a cyclical process [5] with four key components: learners, data, methods, and actions (Fig. 1). A generic LA cycle goes through the phases of *i*) identifying the *learner(s)* and the context in which learning takes place; *ii*) collecting relevant *data*, *iii*) selecting and applying analytics *methods* appropriate for the given learner, learning context, and data, and *iv*) *acting* on the analytics results, often through different forms of pedagogical interventions. This cyclical model bears a lot of resemblance to the CRISP-DM model [6], widely adopted for Data Science (DS) projects. In fact, at the first encounter, LA might be considered as the application of DS in the educational domain. Nonetheless, while the focus on data and computational methods are common to both LA and DS, the two fields differ in some important ways. First, in DS, the primary focus is on the development of high-performance computational models (e.g., prediction models), with less attention to the theoretical

grounding of the model and the ability to explain the phenomena being modelled (e.g., learners at risk of failing a course). On the other hand, LA is focused on supporting evidence-based decision making of distinct participants in the learning process. Therefore, in LA, the development of computational models is first and foremost led by the objective of understanding the learning process. That understanding serves as the basis for acting, that is, taking pedagogical interventions. Henceforth, in LA, model development needs to be grounded in sound pedagogical theory and informed by the specificities of the learning context. Learning context has been recognised as particularly important in model development and results

interpretation [7]-[9]. Likewise, to offer grounds for pedagogically sound interventions, both research questions and methodologies need to be theoretically grounded in well-established learning theories. In short, LA research is not data-driven, as it is often the case in DS, but it makes use of data in a manner shaped by the appropriate learning theory and particularities of the learning context. Furthermore, the importance of understanding a computational model and what can be learnt from it about the learning process and / or learners, is the reason why LA often relies on relatively simple machine learning models, while deep learning models have been rarely adopted.

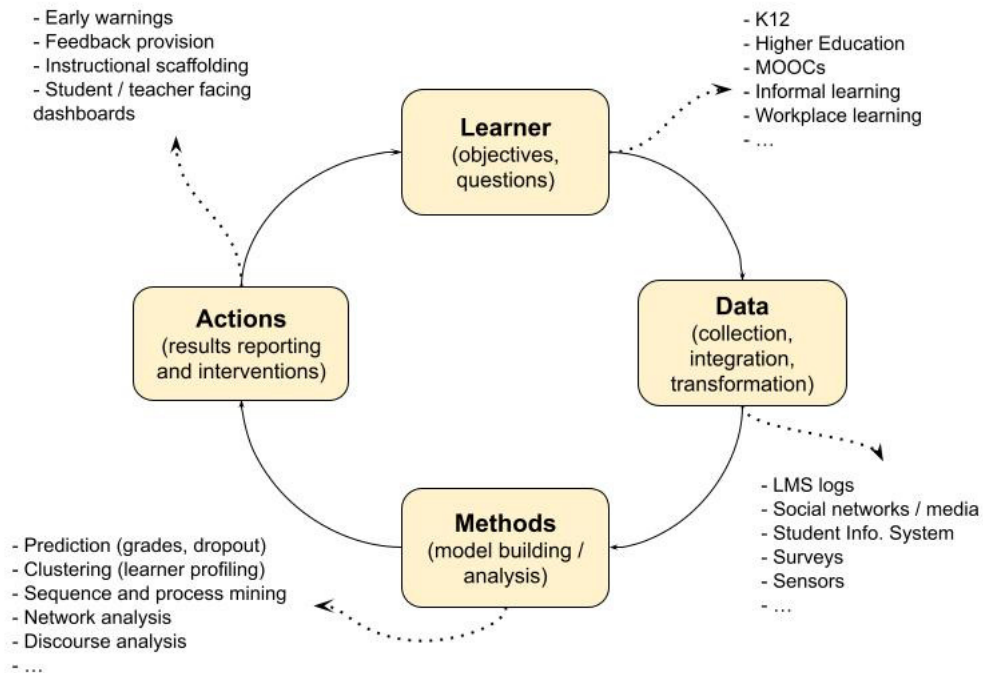


Fig.1 The cyclic Learning Analytics model

Fig. 1 illustrates the cyclical LA model and its key components. Since the overall objective of LA is to understand and optimise learning and the environments in which it occurs, any LA effort starts with identifying *learner(s)* and the learning context to be studied. This ensures that all the subsequent phases of the LA cycle are driven by the objective to support and/or advance learning for the given learner(s) and the given learning context [10]. As AI tools are becoming increasingly present in learning and workplace environments, researchers have started exploring concepts such as hybrid intelligence [11] and hybrid human-AI regulation of learning [12], and some argue for a renewed understanding of the notion of learners, one that integrates the AI dimension [13]. While such altered conceptualization of learners opens interesting research pathways, it goes beyond the scope of the current paper and interested readers are referred to [13] to explore more.

The *data* component refers to the collection, integration, and transformation of data. As already noted, LA relies on data from diverse and often multiple sources, among which the most typically used include learning platforms and tools as well as platforms and tools that may be used for learning (e.g., online social networks and social media); student information system, in case of formal education; various kinds of surveys, often administered before and/or after the studied learning process; sensors such as devices for eye-gaze tracking, position tracking, and video recording of learning [14]. The use of a variety of data, often in a combined manner, allows for comprehensive insights into the learning process. Furthermore, triangulation of data from multiple sources contributes to the trustworthiness of the conclusions derived from the data. However, access to multiple data sources is still a privilege of studies done in controlled settings. In natural

learning settings, learning trace data still remains if not the only, then the dominant data source.

The *methods* component refers to a variety of quantitative and qualitative methods that are used in LA research. The most dominant among LA methods are those based on AI, namely on machine learning and natural language processing. Such methods have been used for predictive modelling (e.g., prediction of students' performance in a course or a study program), learner clustering (e.g., learner profiling based on indicators of engagement with course resources and activities), discourse analysis (e.g., analytics of messages exchanged in online communication channels) [10], [14]. Different kinds of network analysis have been used as well. Social network analysis and epistemological network analysis have been primarily used for developing a better understanding of the structure and content of interactions among actors in the learning process [15], whereas psychological networks have been used for studying both static and dynamic characteristics of learners' psychological states [16]. Process and sequence mining, often combined with advanced statistical modelling, have been used to study the dynamics of learning processes, especially self-regulated learning [14], [17].

Finally, the *actions* component refers primarily to the communication of insights obtained through analytics to relevant stakeholders (learners, teachers, program coordinators, etc) and pedagogical interventions. The communication of analytics results is often done through LA dashboards [18], that is, tools that present LA findings, often in the visual form, in order to support informed decision-making and, in case of learner-facing dashboards, to trigger the desired behavioural change. The action may also take the form of pedagogical interventions, aimed at changing the instructional design and/or offering support to learners (e.g., through different kinds of pedagogical scaffolds) based on the evidence obtained through analytics.

This cyclical model of LA will be used in the next section as the framework for examining the interplay of LA and AI. In particular, by relying on the findings of recent empirical work in LA and related fields of Educational Data Mining and AI in Education, we will explore how key LA components relate to and may benefit from the latest developments in AI, and especially Generative AI.

III. THE INTERPLAY OF LEARNING ANALYTICS AND ARTIFICIAL INTELLIGENCE

This section explores the interplay of LA and AI from the perspective of data, methods, and actions components of the LA cyclic model. For each component, we present how it has been advanced through the use of AI and / or how it has been used to better understand the role / impact of AI on learning. Note that the learner component is not considered due to the

paper's focus on the empirically explored and evidenced interaction of LA and AI, and such efforts, so far, have been based on unaltered notion of learner.

A. Data

Learning traces have been used in a wide variety of LA tasks, most often for predictive modelling and detection of behavioural patterns reflective of the adopted learning tactics and strategies¹. For example, using learning traces from a Coursera course, Jovanovic et al. [16] identified three distinct patterns of learners' interaction with the course activities during individual learning sessions. By considering the visual representation of the identified patterns (Fig. 2) from the perspective of the course design, three learning tactics were identified: assessment-oriented, mastery-oriented, and mixed. These tactics were then used to cluster learners, to identify strategy-based learner profiles. This and similar analysis of learning traces allow LA researchers to understand how learners approach distinct learning and assessment tasks. In other words, analytics of learning traces allow for answering the "what" question – e.g., what learning tactics and strategies a learner has chosen in a course or a module within the course. However, learning traces alone do not allow for answering the "why" questions related to the detected tactics and strategies. These include questions such as why a particular tactic or strategy was selected for the given learning task? Why did a learner switch from one tactic to another and why in a particular moment in time? To answer such questions, data about the learner's internal state (e.g., perceptions, intentions, motivation) are needed.

Some recent studies have also empirically demonstrated the relevance of learner internal factors in predictive modelling. For example, the study presented in [9] analysed a large number of potential predictors of students' academic success (i.e., indicators of academic success derived from trace data), in order to identify predictors that would be relevant across several courses in a study program and thereby, at least partially, enable cross-course portability of predictive models. The study relied on learning traces from a large, homogeneous sample of courses from a healthcare degree program (15 distinct courses, with 50 course offerings). The study results show that behaviour-based indicators explain only a very small percentage of the variability in student achievement, while a significant portion of the variability comes from the students' personal (internal) characteristics. This and similar studies confirmed the intuition about the importance of considering factors characterising learners' internal state when building LA models.

Data about Learner's Internal States. While the relevance of learners' internal factors has been well recognised, collection of data about such factors is still a challenge. Traditionally, such data have been collected through self-reports in the form of surveys, often administered at the beginning and / or

¹ Learning tactic refers to a specific cognitive routine that a learner adopts when solving a particular learning task, whereas a learning strategy is a

specific way the learner selects, applies, and modifies learning tactics when working towards a set learning goal.

at the end of a course. However, such data collection approaches do not allow for capturing the dynamics of learners' motivation, emotions, goal orientations, cognitive load, and other relevant internal factors [19], [20]. In a systematic literature review of LA as a research field, Dawson and colleagues have well recognised challenges associated with learner data collection and noted that “despite the recent advances in multimodal LA, data concerning social and personal dimensions such as motivations, emotions, health and culture are reliant on self-reports or collected from expensive and intrusive equipment” [21]. Current LA research seeks to overcome this challenge through a variety of approaches that all share a common trait, namely the reliance on non-intrusive methods and tools to collect real-time longitudinal data about learners' internal state. As outlined below, some of these approaches rely on human computer interaction to collect data directly from learners, whereas other leverage AI to indirectly obtain (i.e., extract) data about learners' internal states from traces of learner actions and interaction artefacts.

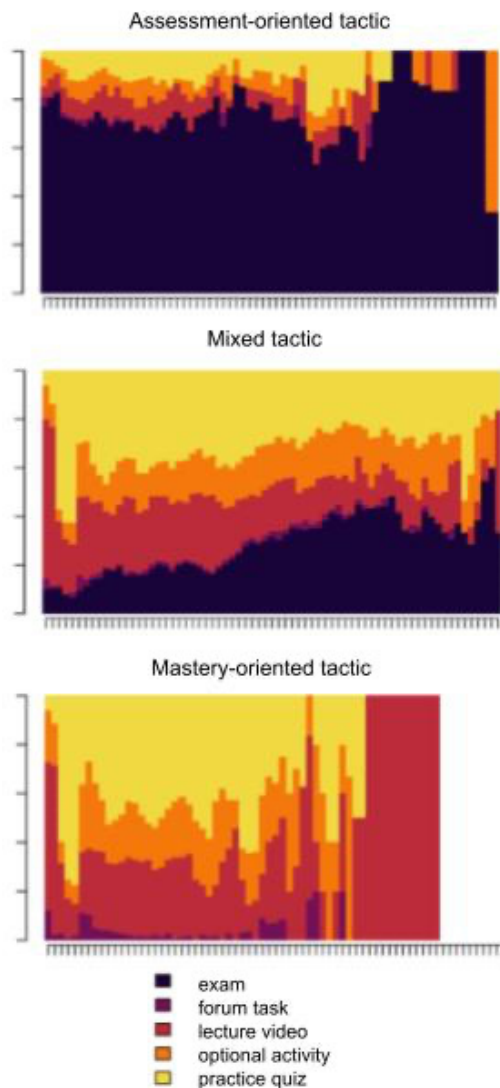


Fig. 2 Illustration of distinct patterns of learners' interaction with course activities, reflective of the adopted learning tactics [16]

Ecological momentary assessment (EMA) is a specific protocol for systematic, longitudinal collection of data about individuals' internal states, by asking a few quick questions in regular time intervals over a longer period of time (e.g., a few weeks or a few months) [22]. Such a protocol is typically operationalised through a mobile phone app, which allows for easy data collection. For example, Fried and colleagues [23] applied EMA to collect data on the psychological and emotional state of students during the first few months of the COVID19 pandemic, and used network analysis to timely identify factors that may cause mental health problems. More recently, Saqr et al [24] used data collected through EMA to build predictive models of distinct self-regulated learning dimensions (e.g., effort regulation, metacognition, motivation and emotions) for individual students. This study also exemplifies an increasing interest of LA research in idiographic analysis, that is, analysis focused on individual students ($N=1$), as it has been shown that conclusions derived from cohort level analysis rarely apply to individual students [25], [26].

As an alternative to direct data collection from learners, LA researchers have also explored the use of AI to automate real-time detection of learners' affective states and emotional engagement, using either trace data alone or traces combined with messages exchanged in online communication channels. An interesting example of the former approach is the work of Hutt et al. [27] who used domain- and platform-independent activity features (e.g., viewing a video lecture, taking a quiz, accessing the discussion board) and state-of-the-art machine learning algorithms to identify 18 distinct emotional states. The authors collected a very large dataset (traces from 69,174 students in 1,898 U.S.A. schools over an entire school year) and built predictive models of learners' affective states that generalised to new students across the two examined domains. However, the models' accuracy was still far from the desired. On the other hand, Liu et al. [28] proposed a state-of-the-art text classification model (BERT-CNN) to identify emotional engagement (positive, negative, confusion) of students in a Massive Open Online Course (MOOC), starting from messages exchanged in the MOOC's discussion forum. While the presented model outperformed alternative models, further improvements are needed before such a model can be used as a trustworthy source of information about learners' emotional engagement in a learning task.

Sophisticated conversational agents, enabled by the latest generation of Generative AI, promise to offer a viable alternative to traditional self-reporting instruments for the collection of data about learners' internal state. By engaging learners in a conversation, instead of presenting them with often long and dull surveys, a chatbot-based data collection approach may prove to be both enticing and effortless and thus increase response rates and quality. To our knowledge, there are still no empirical results on such uses of conversational agents in educational contexts. However, as agent capabilities required for such a task (e.g., proactiveness, goal awareness) are the subject of active research (see e.g., [29]), it is realistic

to expect the use of conversational agents as means of collecting self-reports data in a more natural and enticing manner.

Data Privacy. Data sets used in LA research and practice often contain sensitive data (e.g., student grades, student IP addresses, indicators of psychological and/or physiological state). Data about learners' internal states, discussed above, represent a particularly sensitive category. As the collection of sensitive data increases, concerns regarding privacy protection grow, as well. In general, the continuous increase in volume and diversity of data being collected in educational settings as well as the growing potential for data misuse through the use of advanced technologies (AI included), have made data privacy an area of high concern for LA practice and high relevance and priority for LA research [30].

Traditional approaches to data anonymization have proven insufficient, signalling the need for more robust approaches. For example, it has been shown that unsupervised machine learning techniques can be used to access sensitive student data, despite data anonymization prior to its publishing [31]. An additional challenge is that traditional approaches to protecting data privacy often come at the expense of data utility for LA [32]. In particular, as data privacy increases, data utility, as reflected, for example, in the accuracy of predictions based on that data, declines significantly. All of the above, as well as the general consensus on the need for ethical and responsible use of data in LA, both legally and socially, have led to an increased interest in more robust forms of data anonymization.

Synthetic data represents a state-of-the-art solution for preserving data privacy in highly sensitive domains, such as education, health, and finance [33]. To preserve data privacy, synthetic data, generated by mimicking the characteristics of the original data, is made publicly available instead of the original data. Recent research has demonstrated that, unlike earlier approaches to educational data protection, the use of synthetic data meets the requirements of both data utility and privacy (see, for example, [32], [34]). Sharing of educational data, enabled through the use of synthetic data, is highly important for LA research since it is not unusual that due to the lack of data access, some research objectives need to be abandoned. Furthermore, data sharing is necessary for replication of published research, which is the cornerstone of Open Science.

It is important to mention that until recently, the generation and use of synthetic data was largely limited to structural data, namely tabular data and time series. However, the development of generative AI not only allowed for more sophisticated generation of structured data [33], but also opened opportunities for generating synthetic textual data and multimedia content. For example, to address the problem of limited training data for one-on-one tutoring system, Shan et al. [35] proposed a data augmentation pipeline that leverages Generative AI (GPT-3.5) to create synthetic, multi-labeled dialog data. Similarly, to address limited training data for grounded dialog systems (e.g., tutoring agents), Bao et al. [36] proposed a synthetic data generation framework for grounded dialogues,

which leverages Generative AI (T5) to transform the given dialog flow (i.e., a sequence of knowledge pieces to be covered in a dialog) into a fluent dialog.

Another recent approach to protecting student privacy is to use large language models to identify and remove personally identifying information (PIIs) from messages exchanged in online communication channels. An example is a recent work by Singhal et al. [37] that assessed GPT-4's performance in de-identifying data from discussion forums in nine MOOCs. Overall, the results show high recall (0.958), but low precision (0.526). The tool proved highly successful in identifying PIIs, even identifying cases missed by humans when redacting data. However, it over-redacting names and locations that do not represent PIIs.

B. Methods

From the perspective of LA methods, the interplay of LA and AI comes in two main forms: 1) the use of AI to augment or facilitate LA modelling / methodological approaches, and 2) the use of LA methods to study AI-human interaction in various learning contexts.

Regarding the former aspect of the LA-AI interplay, some methods often used in LA, such as discourse analysis and epistemological network analysis, require qualitative coding of textual content exchanged in learning related interactions (e.g., messages exchanged or comments shared in a collaborative learning task). Qualitative coding has traditionally been a manual task, requiring a lot of time and effort. The latest generation of large language models (LLMs) has provided solid technological grounds for exploring the potentials of semi- or fully automating this task. For example, Hou et al. [38] explored the effectiveness of prompt engineering and fine-tuning approaches for deductive coding of social annotations. In deductive coding, the categories (codes) to be used in the coding task are predefined and often originate either from a relevant theoretical framework or prior empirical research. Categories used for coding can be context dependent or context independent. Context-independent are those categories for which access to individual pieces of content (e.g., a message or a comment) is sufficient to do the coding. On the other hand, context-dependent categories require understanding of the given piece of content in relation to contextually related pieces of content (e.g., previous messages or some external materials), to properly do the coding. In their study, Hou and colleagues [38] considered both kinds of codes and examined the performance of GPT-3.5-turbo adapted to the coding task through prompt engineering or fine tuning. The study results demonstrated that prompt engineering enabled fair to substantial agreement with expert-labelled data across various coding dimensions. Somewhat better results, that is, higher level of agreement, were achieved with fine tuning. As was expected, in both cases, agreement was higher for context-independent than context-dependent categories. In a related study, Barany et al. [39] explored the role that could be played by LLMs, specifically GPT-4, in the process of devel-

oping a codebook for a qualitative coding task, that is, establishing a set of categories to be used for qualitative coding. This is, again, a task that has been done exclusively by researchers. The study compared four approaches to codebook development – a fully manual approach, a fully automated approach, and two approaches that relied on GPT-4 within specific steps of the codebook development process. The study findings suggest that GPT-4 can be valuable for improving qualitative codebooks for use in educational research, but human participation is still essential.

The other form of interchange between LA and AI, namely the use of LA methods to study the interaction of humans and AI in learning situations, is well exemplified in a recent study by Fan et al. [40]. In particular, to examine how students' interaction with Generative AI during an essay revision task compares to interaction with other, more traditional forms of support, Fan and colleagues conducted an experimental study in which they randomly split students into four conditions: one control (no support) and three experimental conditions, each corresponding to a distinct form of support offered during the essay revision task: (human) teacher, ChatGPT, and a checklist suggesting things to focus on when revising the essay. To examine students' interaction with these distinct means of support, the researchers collected learning traces, namely log data, mouse movements, keyboard interaction data, and eye gaze data. The collected traces were parsed into micro-level learning (cognitive and metacognitive) processes which were further analysed through process mining to reveal differences in how interaction with the available help unfolded over the task. This analysis revealed different patterns of interaction with distinct sources of support. In the ChatGPT group, the dominant pattern had a form of back-and-forth between the use of ChatGPT and the very task of revising the essay, whereas other cognitive (e.g., (re-)reading) and metacognitive (e.g., orientation, evaluation) processes were almost absent. On the other hand, the group that interacted with the human teacher did not inhibit, but rather enhanced, connections between essay revising and other learning processes. Furthermore, while the ChatGPT group had significantly higher scores on the revised essays compared to the other conditions, the conditions did not significantly differ in terms of knowledge gain and transfer, nor in the task motivation. Overall, the study findings suggest potential problems of over-reliance on Generative AI and metacognitive laziness, the latter meaning that, when interacting with Generative AI, students tend to leave their metacognitive capacities (monitoring, evaluation, adaptation) dormant. Similar conclusions were reported by Darvishi et al. [41] based on a large randomised controlled study (1625 students across 10 courses). Using LA methods, the study examined if students would learn from regular, detailed, and personalised feedback provided by an LLM-based assistance tool, so that they would be able to exhibit similar behaviour when the assistance is not available. The results showed that students were able to effectively self-regulate their learning with the AI assistance, but with the removal of this support, their performance significantly

dropped. In other words, the students tended to rely on rather than learn from the AI assistance. This and similar findings suggest that with the increasing presence of AI in education, pedagogical interventions that motivate student agency and collaboration with (instead of pure reliance on) AI will be increasingly needed.

A follow-up of the abovementioned study by Fan and colleagues, employed LA methods to examine students' interaction with the human teacher and ChatGPT from the help-seeking perspective [40]. In particular, screen recordings of the students' exchanges with the teacher or ChatGPT were (manually) coded based on the adopted help-seeking theoretical model and the resulting codes served as the input to temporal analysis (process mining) of student - teacher / ChatGPT communication. The resulting process models suggested very different patterns of help-seeking: compared to the human teacher group, in the ChatGPT group, learners asked more "executive" questions (i.e., questions focused on getting direct solutions), and accepted ChatGPT's assistance as is, without evaluation. Furthermore, the students' self-reports after the study revealed lower "social cost" in the ChatGPT group compared to the group working with the human teacher. In other words, students reported being more at ease to seek help from ChatGPT as there were no risks of embarrassing oneself.

Another interesting example of using LA methods to better understand students' use of Generative AI is the study by Brender et al. [42] that examined distinct patterns of student interaction with ChatGPT in the context of a graduate-level robotics course. By clustering students based on the features derived from prompts that students wrote when seeking help from ChatGPT, the researchers identified three profiles (clusters) of ChatGPT use that differed in terms of learning and task performance: i) Debuggers, who requested solutions and error fixes; ii) Conceptual explorers, who sought to understand concepts, tasks, or code, and iii) Practical developers, who exclusively asked for task solutions. While Debuggers had the best task performance, like Practical developers, they were less likely to translate performance into conceptual understanding. On the other hand, Conceptual Explorers had better overall learning outcomes compared to the other two profiles. This study offers yet another confirmation that over-reliance on Generative AI, while often beneficial for short-term performance goals, may inhibit a true mastery of new knowledge and skills.

LA researchers are also experimenting with AI-based pedagogical interventions that include student interaction with more than one AI-based agent. For example, an ongoing study in the domain of medical education enrolls two LLMs in the task of helping student doctors to learn how to talk to a patient [40]. In particular, one LLM is acting as a patient, whereas the other takes on the role of a senior medical doctor "observing" the interaction between the "patient" and the student doctor and providing feedback to the student.

Finally, there are some nascent approaches to using LA to assess human-AI collaborative work. These have been motivated by the recognition that AI systems and tools are becoming an intrinsic part of various kinds of professions and that the future of work would include different forms of human-AI collaboration. Hence, it will be the task of education to help learners develop knowledge and skills required for a thriving human-AI collaboration as well as to assess such collaboration. The assessment of human-AI collaboration includes not only evaluation of the outcome of a collaborative task, but also evaluation of the processes that led to those outcomes [43]. An example of this line of research is a recent work by Cheng et al. [44] that proposed a LA-based method for assessing collaborative writing of humans and Generative AI. The method relies on learning trace data collection, their mapping to learning processes, and finally epistemological network analysis of student-AI exchanges.

C. Actions

Learning analytics dashboards are a primary method for delivering analytics results to end users, thus facilitating evidence-based decision-making and actions. However, a persistent challenge has been communicating LA feedback in a way that end users, who may lack technical expertise, can accurately interpret and act upon [18]. To address this challenge, researchers have explored ways for augmenting LA dashboards with Generative AI. For example, Yan et al. [45] proposed VizChat, an open-sourced, prototype chatbot designed to augment LA dashboards with contextualised, AI-generated explanations of visually presented LA results. The objective is to improve user comprehension of the dashboard without overwhelming the user with excessive information. To that end, VizChat leverages multimodal Generative AI (GPT-4V) and Retrieval Augmented Generation (RAG) to offer on-demand, contextually relevant explanations of specific visualisations as well as a summary of information integrated from multiple visual depictions of LA results. To increase the transparency and contribute to trust in the feedback communicated through the dashboard, the tool also offers detailed information about the data sources used and analytics processes behind each visualisation. Still, the informativeness and usability of VizChat has yet to be verified through more comprehensive empirical studies with students and teachers.

Another interesting approach to advancing communication of LA feedback is storytelling augmented with Generative AI. The use of storytelling either as an alternative or a complement to LA dashboards has already been explored (e.g., [46]), especially in the context of multimodal LA, where, due to the use of multiple data sources, the challenge of clear results communication is especially high [47]. Aiming to further facilitate communication of LA feedback to learners and make it more appealing, Milesi and colleagues [48] explored the combined use of Generative AI and data comics, the latter being an emergent storytelling format for helping end users

(non-expert) understand complex data and analytics. In particular, the researchers used MidJourney, an image generation AI tool, and a graphics illustration tool to create personal data comics about students' multimodal LA data. The initial evaluation of this approach with nursing students showed that while students found Generative-AI-augmented data comics appealing and enjoyable, they also expressed concerns that such a form of communicating insights from data lack the professionalism required for the given learning context (professional education). While probably not suitable for adult learners, this approach holds promise for young learners.

A well-recognized limitation of LA dashboards is the unidirectional communication of LA data and feedback. Research on educational feedback has shown that such (one-way) communication of feedback is far from optimal [49]. What is preferable is a dialog form, that is, bidirectional communication that allows for better dealing with any potential problem revealed through analytics or resolving any potential misinterpretation of the originally communicated analytics findings. Conversational chatbots, enabled by Generative AI, have opened opportunities for engaging students in such dialogic feedback. Rich literature on pedagogical agents [50], which predates the recent Generative AI developments, may offer strong foundation for such conversational agents. However, at the point of writing this manuscript, empirical findings that may confirm the expected benefits of Generative AI for dialogic feedback provision are still lacking.

Some recent studies have examined the use of the latest generation of LLMs for automated generation of feedback on student produced content, with the ultimate objective of helping students improve their writing. For example, Hutt et al. [51] examined the use of ChatGPT for providing students with feedback on peer feedback, that is, helping students learn what constitutes "good" feedback and how to provide it². To understand the potentials of the latest generation of LLMs compared to earlier AI-based solutions, Hutt and colleagues compared ChatGPT with traditional text classification models in estimating the quality of peer feedback, according to the given rubric. The traditional AI models proved more accurate, while the advantage of ChatGPT was that it produced explanations of the assigned quality category.

IV. CONCLUSION

This paper explored the interplay between Learning Analytics (LA) and Artificial Intelligence (AI), as evidenced in recent LA research. It highlighted both the benefits AI has brought to LA and the ways in which LA has been used to enhance our understanding of AI's role and impact on learning, particularly with Generative AI. All this suggests that LA community has made significant contributions both in:

- using AI to address long-standing challenges in LA research, such as ensuring data privacy and advancing LA dashboards.

²Peer feedback is considered a powerful learning strategy as it offers learning opportunities both for the learner receiving feedback and the learner

providing feedback. However, students often lack knowledge regarding what constitutes "good" feedback and need to learn how to provide it.

- using LA to gain insights into learners' interactions with AI, such as identifying learners' tendency to over-rely on AI and neglect metacognitive processes.

AI introduces new opportunities and challenges for LA research while also equipping researchers with more advanced methods, richer data, and improved ways of communicating analytics results, such as explanations and dialogic feedback. The dynamic between LA and AI promises to continually yield relevant insights into the evolving role of AI in learning. For these insights to be effectively integrated into educational practice, the active engagement of all stakeholders is crucial, alongside public policies that recognize the importance of timely, evidence-based decision-making in the era of Generative AI.

REFERENCES

- [1] C. Lang, A. F. Wise, A. Merceron, D. Gašević, and G. Siemens, "What is Learning Analytics?," in *The Handbook of Learning Analytics*, 2nd ed., Vancouver, Canada: SoLAR, 2022, pp. 8–18. [Online]. Available: <https://www.solaresearch.org/publications/hla-22/hla22-chapter1/>
- [2] H. Ouhachi, D. Spikol, and B. Vogel, "Research trends in multimodal learning analytics: A systematic mapping study," *Computers and Education: Artificial Intelligence*, vol. 4, p. 100136, Jan. 2023, <http://dx.doi.org/10.1016/j.caeai.2023.100136>
- [3] P. H. Winne and A. E. Hadwin, "Studying as Self-Regulated Learning," in *Metacognition in Educational Theory and Practice*, Routledge, 1998.
- [4] P. H. Winne, "Cognition and Metacognition within Self-Regulated Learning," in *Handbook of Self-Regulation of Learning and Performance*, 2nd ed., Routledge, 2017.
- [5] D. Clow, "The learning analytics cycle: closing the loop effectively," in *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*. New York, NY, USA: ACM, Apr. 2012, pp. 134–138. <http://dx.doi.org/10.1145/2330601.2330636>
- [6] N. Hotz, "What is CRISP DM?," Data Science Process Alliance. Accessed: Jul. 18, 2024. [Online]. Available: <https://www.datasciencepm.com/crisp-dm-2/>
- [7] D. Gašević, S. Dawson, T. Rogers, and D. Gasevic, "Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting academic success," *The Internet and Higher Education*, vol. 28, pp. 68–84, Jan. 2016, <http://dx.doi.org/10.1016/j.iheduc.2015.10.002>
- [8] R. Conijn, C. Snijders, A. Kleingeld, and U. Matzat, "Predicting Student Performance from LMS Data: A Comparison of 17 Blended Courses Using Moodle LMS," *IEEE Transactions on Learning Technologies*, vol. 10, no. 1, pp. 17–29, Jan. 2017, <http://dx.doi.org/10.1109/TLT.2016.2616312>
- [9] J. Jovanović, M. Saqr, S. Joksimović, and D. Gašević, "Students matter the most in learning analytics: the effects of internal and instructional conditions in predicting academic success," *Computers & Education*, vol. 172, p. 104251, Oct. 2021, <http://dx.doi.org/10.1016/j.compedu.2021.104251>
- [10] D. Gašević, S. Dawson, and G. Siemens, "Let's not forget: Learning analytics are about learning," *TECHTRENDS TECH TRENDS*, vol. 59, no. 1, pp. 64–71, Dec. 2014, <http://dx.doi.org/10.1007/s11528-014-0822-x>
- [11] S. Jarvela *et al.*, "Hybrid Intelligence – Human-AI Co-Evolution and Learning in Multirealities (HI)," in *HHAI 2023: Augmenting Human Intellect*, IOS Press, 2023, pp. 392–394. <http://dx.doi.org/10.3233/FAIA230107>
- [12] I. Molenaar, "The concept of hybrid human-AI regulation: Exemplifying how to support young learners' self-regulated learning," *Computers and Education: Artificial Intelligence*, vol. 3, p. 100070, Jan. 2022, <http://dx.doi.org/10.1016/j.caeai.2022.100070>
- [13] L. Yan, R. Martinez-Maldonado, and D. Gašević, "Generative Artificial Intelligence in Learning Analytics: Contextualising Opportunities and Challenges through the Learning Analytics Cycle," in *Proceedings of the 14th Learning Analytics and Knowledge Conference*, Mar. 2024, pp. 101–111. <http://dx.doi.org/10.1145/3636555.3636856>
- [14] A. Palanci, R. M. Yilmaz, and Z. Turan, "Learning analytics in distance education: A systematic review study," *Educ Inf Technol*, May 2024, <http://dx.doi.org/10.1007/s10639-024-12737-5>
- [15] B. Chen and O. Poquet, "Networks in Learning Analytics: Where Theory, Methodology, and Practice Intersect," *Journal of Learning Analytics*, vol. 9, no. 1, Art. no. 1, Mar. 2022, <http://dx.doi.org/10.18608/jla.2022.7697>
- [16] J. Jovanovic, D. Gašević, L. Yan, G. Baker, A. Murray, and D. Gasevic, "Explaining trace-based learner profiles with self-reports: The added value of psychological networks," *Journal of Computer Assisted Learning*, 2024, <http://dx.doi.org/10.1111/jcal.12968>
- [17] J. Du, K. F. Hew, and L. Liu, "What can online traces tell us about students' self-regulated learning? A systematic review of online trace data analysis," *Computers & Education*, vol. 201, p. 104828, Aug. 2023, <http://dx.doi.org/10.1016/j.compedu.2023.104828>
- [18] R. Kalliisa, K. Misiejuk, S. López-Pernas, M. Khalil, and M. Saqr, "Have Learning Analytics Dashboards Lived Up to the Hype? A Systematic Review of Impact on Students' Achievement, Motivation, Participation and Attitude," in *Proceedings of the 14th Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2024, pp. 295–304. <http://dx.doi.org/10.1145/3636555.3636884>
- [19] M. Zhou and P. H. Winne, "Modeling academic achievement by self-reported versus traced goal orientation," *Learning and Instruction*, vol. 22, no. 6, pp. 413–419, Dec. 2012, <http://dx.doi.org/10.1016/j.learninstruc.2012.03.004>
- [20] D. Gasevic, J. Jovanovic, A. Pardo, and S. Dawson, "Detecting learning strategies with analytics: links with self-reported measures and academic performance," *Journal of Learning Analytics*, vol. 4, no. 2, pp. 113–128, Jul. 2017.
- [21] S. Dawson, S. Joksimovic, O. Poquet, and G. Siemens, "Increasing the Impact of Learning Analytics," in *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*. New York, NY, USA: ACM, Mar. 2019, pp. 446–455. <http://dx.doi.org/10.1145/3303772.3303784>
- [22] S. Zirkel, J. A. Garcia, and M. C. Murphy, "Experience-Sampling Research Methods and Their Potential for Education Research," *Educational Researcher*, vol. 44, no. 1, pp. 7–16, Jan. 2015, <http://dx.doi.org/10.3102/0013189X14566879>
- [23] E. I. Fried, F. Papanikolaou, and S. Epskamp, "Mental Health and Social Contact During the COVID-19 Pandemic: An Ecological Momentary Assessment Study," *Clinical Psychological Science*, vol. 10, no. 2, pp. 340–354, Mar. 2022, <http://dx.doi.org/10.1177/21677026211017839>
- [24] M. Saqr, R. Cheng, S. López-Pernas, and E. D. Beck, "Idiographic artificial intelligence to explain students' self-regulation: Toward precision education," *Learning and Individual Differences*, vol. 114, p. 102499, Aug. 2024, <http://dx.doi.org/10.1016/j.lindif.2024.102499>
- [25] O. Poquet, K. Kitto, J. Jovanovic, S. Dawson, G. Siemens, and L. Markauskaite, "Transitions through lifelong learning: Implications for learning analytics," *Computers and Education: Artificial Intelligence*, vol. 2, p. 100039, Nov. 2021, <http://dx.doi.org/10.1016/j.caeai.2021.100039>
- [26] M. Saqr, "Group-level analysis of engagement poorly reflects individual students' processes: Why we need idiographic learning analytics," *Computers in Human Behavior*, p. 107991, Oct. 2023, <http://dx.doi.org/10.1016/j.chb.2023.107991>
- [27] S. Hutt, J. F. Grafsgaard, and S. K. D'Mello, "Time to Scale: Generalizable Affect Detection for Tens of Thousands of Students across An Entire School Year," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2019, pp. 1–14. <http://dx.doi.org/10.1145/3290605.3300726>
- [28] S. Liu, S. Liu, Z. Liu, X. Peng, and Z. Yang, "Automated detection of emotional and cognitive engagement in MOOC discussions to predict learning achievement," *Computers and Education*, vol. 181, no. C, May 2022, <http://dx.doi.org/10.1016/j.compedu.2022.104461>
- [29] L. Liao, G. H. Yang, and C. Shah, "Proactive Conversational Agents in the Post-ChatGPT World," in *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York, NY, USA: ACM, Jul. 2023, pp. 3452–3455. <http://dx.doi.org/10.1145/3539618.3594250>
- [30] S. Joksimović, R. Marshall, T. Rakotoarivelo, D. Ladjal, C. Zhan, and A. Pardo, "Privacy-Driven Learning Analytics," in *Manage Your Own Learning Analytics: Implement a Rasch Modelling Approach*, E.

- McKay, Ed., Springer International Publishing, 2022, pp. 1–22. http://dx.doi.org/10.1007/978-3-030-86316-6_1
- [31] E. Yacobson, O. Fuhrman, S. Hershkovitz, and G. Alexandron, “De-identification is Insufficient to Protect Student Privacy, or – What Can a Field Trip Reveal?,” *Journal of Learning Analytics*, vol. 8, no. 2, Art. no. 2, Sep. 2021, <http://dx.doi.org/10.18608/jla.2021.7353>.
- [32] Q. Liu, M. Khalil, J. Jovanovic, and R. Shakya, “Scaling While Privacy Preserving: A Comprehensive Synthetic Tabular Data Generation and Evaluation in Learning Analytics,” in *Proceedings of the 14th Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2024, pp. 620–631. <http://dx.doi.org/10.1145/3636555.3636921>.
- [33] J. Jordan *et al.*, “Synthetic Data -- what, why and how?,” May 06, 2022, *arXiv: arXiv:2205.03257*. <http://dx.doi.org/10.48550/arXiv.2205.03257>.
- [34] C. Zhan, S. Joksimović, D. Ladjal, T. Rakotoarivelo, R. Marshall, and A. Pardo, “Preserving Both Privacy and Utility in Learning Analytics,” *IEEE Transactions on Learning Technologies*, vol. 17, pp. 1655–1667, 2024, <http://dx.doi.org/10.1109/TLT.2024.3393766>.
- [35] D. Shan, D. Wang, C. Zhang, B. Kao, and C. K. K. Chan, “Annotating Educational Dialog Act with Data Augmentation in Online One-on-One Tutoring,” in *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*, Cham: Springer Nature Switzerland, 2023, pp. 472–477. http://dx.doi.org/10.1007/978-3-031-36336-8_73.
- [36] J. Bao *et al.*, “A Synthetic Data Generation Framework for Grounded Dialogues,” in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 10866–10882. <http://dx.doi.org/10.18653/v1/2023.acl-long.608>.
- [37] S. Singhal, A. F. Zambrano, M. Pankiewicz, X. Liu, C. Porter, and R. S. Baker, “De-Identifying Student Personally Identifying Information with GPT-4,” in the *Proceedings of the 17th International Conference on Educational Data Mining*, 2024, pp. 559–565. <http://dx.doi.org/10.5281/zenodo.12729884>.
- [38] C. Hou *et al.*, “Prompt-based and Fine-tuned GPT Models for Context-Dependent and -Independent Deductive Coding in Social Annotation,” in *Proceedings of the 14th Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2024, pp. 518–528. <http://dx.doi.org/10.1145/3636555.3636910>.
- [39] A. Barany *et al.*, “ChatGPT for Education Research: Exploring the Potential of Large Language Models for Qualitative Codebook Development,” in *Artificial Intelligence in Education*, Cham: Springer Nature Switzerland, 2024, pp. 134–149. http://dx.doi.org/10.1007/978-3-031-64299-9_10.
- [40] Y. Fan, “Learning and Regulating with ChatGPT: What an Experimental Study Tells Us,” Jun. 25, 2024. Accessed: Jul. 18, 2024. [Online]. Available: <https://www.youtube.com/watch?v=YdWePjSYrzw>
- [41] A. Darvishi, H. Khosravi, S. Sadiq, D. Gašević, and G. Siemens, “Impact of AI assistance on student agency,” *Computers & Education*, p. 104967, Nov. 2023, <http://dx.doi.org/10.1016/j.compedu.2023.104967>.
- [42] J. Brender, L. El-Hamamsy, F. Mondada, and E. Bumbacher, “Who’s Helping Who? When Students Use ChatGPT to Engage in Practice Lab Sessions,” in *Artificial Intelligence in Education*, Cham: Springer Nature Switzerland, 2024, pp. 235–249. http://dx.doi.org/10.1007/978-3-031-64302-6_17.
- [43] D. Gasevic, “Reimagining Assessment in the Age of Artificial Intelligence,” May 13, 2024. Accessed: Jul. 18, 2024. [Online]. Available: <https://www.youtube.com/watch?v=2Ea3oJysD6s>
- [44] Y. Cheng, K. Lyons, G. Chen, D. Gašević, and Z. Swiecki, “Evidence-centered Assessment for Writing with Generative AI,” in *Proceedings of the 14th Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2024, pp. 178–188. <http://dx.doi.org/10.1145/3636555.3636866>.
- [45] L. Yan *et al.*, “VizChat: Enhancing Learning Analytics Dashboards with Contextualised Explanations Using Multimodal Generative AI Chatbots,” in *Artificial Intelligence in Education*, Cham: Springer Nature Switzerland, 2024, pp. 180–193. http://dx.doi.org/10.1007/978-3-031-64299-9_13.
- [46] G. M. Fernandez Nieto, K. Kitto, S. Buckingham Shum, and R. Martinez-Maldonado, “Beyond the Learning Analytics Dashboard: Alternative Ways to Communicate Student Data Insights Combining Visualisation, Narrative and Storytelling,” in *LAK22: 12th International Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2022, pp. 219–229. <http://dx.doi.org/10.1145/3506860.3506895>.
- [47] R. Martinez-Maldonado, V. Echeverria, G. Fernandez Nieto, and S. Buckingham Shum, “From Data to Insights: A Layered Storytelling Approach for Multimodal Learning Analytics,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2020, pp. 1–15. <http://dx.doi.org/10.1145/3313831.3376148>.
- [48] M. E. Milesi *et al.*, “‘It’s Really Enjoyable to See Me Solve the Problem like a Hero’: GenAI-enhanced Data Comics as a Learning Analytics Tool,” in *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2024, pp. 1–7. <http://dx.doi.org/10.1145/3613905.3651111>.
- [49] B. Maheshi, W. Dai, R. Martinez-Maldonado, and Y.-S. Tsai, “Dialogic feedback at scale: Recommendations for learning analytics design,” *Journal of Computer Assisted Learning*, 2024, <http://dx.doi.org/10.1111/jcal.13034>.
- [50] P. Sikström, C. Valentini, A. Sivunen, and T. Kärkkäinen, “How pedagogical agents communicate with students: A two-phase systematic review,” *Computers & Education*, vol. 188, p. 104564, Oct. 2022, <http://dx.doi.org/10.1016/j.compedu.2022.104564>.
- [51] S. Hutt *et al.*, “Feedback on Feedback: Comparing Classic Natural Language Processing and Generative AI to Evaluate Peer Feedback,” in *Proceedings of the 14th Learning Analytics and Knowledge Conference*. New York, NY, USA: ACM, Mar. 2024, pp. 55–65. <http://dx.doi.org/10.1145/3636555.3636850>.

Attentiveness on criticisms and definition about Explainable Artificial Intelligence

Francisco Herrera

Dept. of Computer Science and Artificial Intelligence
Andalusian Institute of Data Science and Computational Intelligence (DaSCI)
University of Granada, Spain.
Email: herrera@decsai.ugr.es

Abstract—The emergence of deep learning at the beginning of the last decade has driven the advancement of complex models, culminating in the development of large language models and generative AI. These models represent the summit of size and complexity. Explainability should be an option that plays a key role in enabling understandable the AI-assisted decision-making and ensuring accountability. This contribution delves into the complexities of explainable artificial intelligence (XAI) through various perspectives, considering the extensive and growing body of literature. Our discussion begins by addressing the challenges posed by complex data, models, and high-risk scenarios. Given the rapid growth of the field, it is essential to tackle the criticisms and challenges that emerge as it matures, requiring thorough exploration. This contribution explores them, along with three aspects that may shed light on them. First, it is focused on the lack of definitional cohesion, examining how and why is defined XAI from the perspectives of audience and understanding. Second, it explores XAI explanations, bridging the gap between complex AI models and human understanding. Third, it is crucial to consider how to analyze and evaluate the maturity level of explainability, from a triple dimension, practicality, governance and auditability.

Index Terms—eXplainable Artificial Intelligence, explanations, metrics, audience.

I. INTRODUCTION

IN RECENT years, the rapid advancement of artificial intelligence (AI) has led to the development of increasingly complex models capable of performing tasks with remarkable accuracy. However, the opacity of these models, often referred to as "black-box AI" [1], has raised significant concerns regarding their interpretability and trustworthiness. We work with complex data, complex black box models and complex scenarios dealing with high risks problems. Explainable AI (XAI) has emerged as a critical field of research aimed at addressing these concerns by providing transparent and understandable explanations for the AI-assisted decision-making. AI-assisted decision-making refers to the process where AI systems provide recommendations or insights to help humans make decisions.

The European Union greenlit the first major AI law, AI Act¹ in december 2023, approved in march 2024, and published on july 2024. It will regulate the development, use, and application of AI. Its goal is to ensure AI systems used and

developed in the EU are safe and trustworthy. "The adoption of the AI Act is a significant milestone for the European Union. This landmark law, the first of its kind in the world, addresses a global technological challenge that also creates opportunities for our societies and economies. With the AI Act, Europe emphasizes the importance of trust, transparency and accountability when dealing with new technologies while at the same time ensuring this fast-changing technology can flourish and boost European innovation." said recently on the occasion of the approval Mathieu Michel, Belgian secretary of state for digitisation, administrative simplification, privacy protection, and the building regulation².

Explainability is both in the fundamental principles associated with the European trustworthy AI definition³ (respect for human autonomy, prevention of harm, equity, and explainability), and in UNESCO's ethical principle⁴, number 7; Transparency and explainability. It is as well as being part of the European transparency requirement for high risk problems: "The behavior of AI systems must be able to be monitored or traced, or in other words, record all their procedures, from the data acquisition and annotation process, to each of the decisions they make. It is therefore vital that AI systems are explainable, in order to understand the decisions they make based on certain input data. It is clear that making AI processes and decisions explainable is essential." In January 2023, the National Institute of Standards and Technology (NIST) published the Artificial Intelligence Risk Management Framework (AI RMF 1.0)⁵, which includes a similar list of trustworthy AI characteristics (it uses the term characteristic with a similar meaning to requirement), highlighting characteristics such as "safety and resilience", "explainability and interpretability" (separate from "transparency").

This is a general scenario under which we analyze the usability, utility and future of AI. Therefore, XAI is recognized as a crucial area with significant potential to foster trust, en-

²<https://www.consilium.europa.eu/en/press/press-releases/2024/05/21/artificial-intelligence-ai-act-council-gives-final-green-light-to-the-first-worldwide-rules-on-ai/>

³Ethics Guidelines for Trustworthy Artificial Intelligence. HLEG-AI, 2019 <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

⁴Recommendation on the Ethics of Artificial Intelligence, <https://unesdoc.unesco.org/ark:/48223/pf0000381137>

⁵AIRMF-NIST. Artificial Intelligence Risk Management Framework (AI RMF 1.0). <https://doi.org/10.6028/NIST.AI.100-1>

This work was supported by the national project PID2023-150070NB-I00, Ministry of Science, Innovation and Universities

¹AI Act, <https://artificialintelligenceact.eu/the-act/>

sure accountability, and enable informed AI-assisted decision-making across various high-risk domains [2], [3], including healthcare, finance, and autonomous systems, public services, among others.

We can read the vast literature on XAI, from the current state of maturity [4] to its challenges [5] and highlighted criticisms [6], [7], [8]. The scientific literature is very prolific, I don't know if too much, it sheds a lot of light, many results, but also many questions. It is essential to approach XAI context with nuance and conduct in-depth analysis to ensure progress is made in the right direction. This exercise aims to briefly address various aspects and questions by analyzing authors' opinions, criticisms, established working areas, and XAI evaluation.

This contribution explores the complexities of XAI by analyzing various discussions within the literature. It begins by addressing the challenges posed by complex data, models, and high-risk scenarios. As the field matures, it is crucial to thoroughly examine the criticisms and challenges that arise. The paper then focuses on the lack of definitional cohesion, emphasizing the importance of defining XAI from the perspectives of audience and understanding, using an existing definition. We explore XAI techniques and discuss the crucial element of XAI explanations. Finally, we highlight the importance of evaluating the maturity level of explainability and how to measure it, from a triple dimension, practicality, governance and auditability.

The contribution is organized into sections based on the mentioned studies. It features concluding remarks on long way to go to enhance the usefulness of XAI, and also mentioning some topics that have been left untouched or barely explored.

II. COMPLEX DATA, COMPLEX MODELS AND HIGH RISK SCENARIOS

The emergence of deep learning at the beginning of the last decade led us to begin the advance in complex models, up to the large language models and generative AI as the summation of size and complexity. Also during these years, proposals have been made for deep structure neural networks models with different architectures that process various types of data, such as images, video, time series, text and multimodal data. Feature engineering is an essential methodology when working with tabular data and well-structured data, but it falls short when working with the complexity of the data mentioned above. It is not feasible to associate features with images, for example.

This puts us in a context of increasing complexity, which makes us understand less how AI models work, leading non-experts into the abyss of "Without understanding AI, observing the magic of AI". Therefore the desire for explainability becomes a universally accepted goal.

On the other hand, Europe has established the first law on AI, the AI Act⁶, published on 13 July 2024. AI deployment will be graded on a risk-based scale. Technologies with an

unacceptable risk of causing direct harm will be banned. Where AI impacts fundamental human rights or critical systems such as essential infrastructure, public transport, healthcare or wellbeing, it will be classified as "high risk" and subject to increased levels of oversight and accountability. This regulation describes the concept of high-risk based AI systems, as those AI systems that are used in any of the following eight high-risk scenarios:

- biometric identification and categorisation of natural persons,
- management and operation of essential infrastructure,
- education and vocational training,
- employment, management of workers and access to self-employment,
- access to and enjoyment of essential public and private services and their benefits,
- matters related to law enforcement,
- management of migration, asylum and border control, or
- administration of justice and democratic processes.

In a few lines we have shown a global context that has been consolidated in recent years, and where it is necessary to advance in trustworthy AI technologies for the design of responsible AI systems [9], and XAI is a cornerstone.

III. EXAMINING THE TROUBLES AND CRITICISMS IN XAI

Given the maturity that XAI is gaining, reflections are raised on the path it follows and the associated problems. There are works in the literature that criticize XAI for various reasons, criticisms of its relevance in the current context.

Among them, we find specific criticisms about a concrete question as the use of certain XAI measures, for example in [10] is presented arguments demonstrating that Shapley values for explainability can produce misleading information regarding relative feature important. The authors state emphatically in [11] that: "The continued practical use of tools that approximate SHAP scores should be a reason of concern in high-risk and safety-critical domains".

On the other hand, we find deep discussions on the XAI troubles. This contribution focuses on three articles with deep and thoughtful criticisms by their authors [6], [7], [8]. They will be developed below, to end with a brief position on these criticisms.

The paper [6], with the striking title "Dear XAI Community, We Need to Talk!", highlights and discusses eight misconceptions in XAI research. Authors argue on the lack of solid grounds due to:

- *"Proposals for new interpretation techniques that serve no clear purpose;*
- *anecdotal evidence from intuitive-looking heatmaps or "benchmarks" on seemingly relevant criteria are used as a substitute for a clear motivation;*
- *explanations are generated that mislead humans into trusting ML models without the models being trustworthy".*

⁶AI Act, <https://artificialintelligenceact.eu/the-act/>

The misconceptions are collected under the following titles:

- 1) *“Explanation Methods are Purpose-Free”*. A that explanation techniques in XAI should serve at least one practical purpose. Authors emphasize the importance of clearly demonstrating how an explanation technique fulfills its intended purpose. Without a widely accepted definition of explainability or interpretability, the purpose is the key to connecting these techniques to real-world applications. Techniques that lack practical motivation should be viewed with skepticism. If an explanation cannot be shown to help its intended audience, it is likely not useful and should be discarded.
 - 2) *“One Explanation Technique to Rule Them All”*. Authors argue that XAI community members believe that a single best explanation technique, like SHAP, can provide perfect understanding for all purposes. Authors emphasize that the goals of explanations in XAI are diverse, such as auditing models, understanding phenomena, debugging, or enabling users to contest decisions. Different goals require different techniques and hyperparameters. They use counterfactual explanations to illustrate conflicts and trade-offs. For example, age may be excluded in counterfactuals for recourse but included for contesting decisions. But, counterfactuals may not be suitable for understanding the model as they offer limited insights.
 - 3) *“Benchmarks do not Need a Ground-Truth”*. Authors discuss the challenges of benchmarking in XAI. While benchmarks have been successful in ML due to the presence of a ground truth, XAI lacks this central element, making objective comparisons difficult. Authors suggest two ways to progress in XAI: either abandon benchmarks and focus on qualitative evaluation or define benchmarks based on the explanation’s purpose. However, some in the XAI community have taken a less rigorous approach, optimizing explanations for specific properties without clear motivation. This undermines the validity of benchmarks, turning them into promotional tools rather than objective standards.
 - 4) *“We Should Give People Explanations They Find Intuitive”*. Authors criticize the practice of tailoring explanations in XAI to fit human intuition, which may not be faithful to the actual model. They argue that XAI should aim to make the model’s mechanisms transparent rather than convincing people to trust the system. They distinguish between explanations (actual reasons for decisions) and justifications (good reasons for decisions), noting that they often diverge in XAI. They emphasize the need for explanations that are faithful to the causal decision-making process, rather than those designed to be compelling or intuitive.
 - 5) *“Current Deep Nets Accidentally Learn Human Concepts”*. Authors challenge the assumption that deep neural networks learn the same concepts as humans. They argue that while early layers may learn low-level concepts and later layers high-level concepts, this does not mean the model’s reasoning aligns with human logic. They highlight several issues, such as the distributed representation enforced by regularization techniques and the limited impact of manipulating specific neurons. They also point out that effective communication, a key reason for shared human concepts, is not a constraint in ML training. They conclude that techniques like activation maximization may produce misleading results, and it is doubtful that humans will ever fully understand the concepts used by ML models.
 - 6) *“Every XAI Paper Needs Human Studies”*. Authors emphasize the importance of human studies in evaluating explanations in XAI. They highlight two key questions: what conceptually counts as an explanation for a phenomenon, and which explanations are good for specific explainees. While human studies are essential for the latter, the former can be addressed through conceptual analysis and formal tools. Conceptual definitions help narrow down the vast space of possible explanations, guiding the search for good ones. Not all XAI purposes require human studies; for example, formal evaluations can be justified if human studies have already been conducted for that type of explanation.
 - 7) *“XAI Methods can be Wrong”*. Authors discuss the limitations and challenges of saliency-based and model-agnostic explanation techniques like SHAP, LIME, and counterfactuals in XAI. They highlight that while these techniques can be manipulated to provide desired explanations, this does not necessarily mean they are wrong. Instead, they underscore the need for diverse XAI techniques, each illuminating different aspects of a model. They emphasize the importance of developing XAI techniques at various levels of abstraction to provide a comprehensive understanding of model behavior and address real-world purposes.
 - 8) *“Extrapolating to Stay True to the Model”*. Authors discuss how most XAI techniques probe ML models, often in areas where the model has not seen any data, leading to extrapolation. Techniques like LIME, SHAP, and counterfactuals rely on probing the model, but ML models are generally poor at extrapolating to unseen instances. Explanations based on extrapolation may not be reliable. They argue that the explanations should focus on areas where the model is qualified, as probing outside the data manifold makes interpretation blurry and problematic. They emphasize the need for XAI techniques that provide insights within the data manifold for most purposes.
- The paper is accompanied by four steps forward to take (section 4), sharing authors thoughts and intuitions about how they think the field should evolve to become a more substantive discipline. Their steps forward are:
- *Go from purpose to benchmark,*
 - *Be clear what you need to explain and by what,*

- Give clear instructions for how to interpret explanation Techniques, and
- XAI needs interdisciplinarity and expertise.

We must consider this paper as a fairly in-depth analysis of the problems in XAI.

John Zerilli raises an interesting reflection from a philosophical prism in [7]: "XAI has been forced to prioritise interpretability at the expense of completeness, and even realism, so that its explanations are frequently interpretable without being underpinned by more comprehensive explanations faithful to the way a network computes its predictions. While this has been taken to be a shortcoming of the field of XAI, I argue that it is broadly the right approach to the problem." He concludes "for deeper and more comprehensive explanations of automated decisions is urgent, as in some cases it may be, we should naturally expect them, in whatever form is considered practicable by the standards of XAI. But where no such necessity arises, a satisficing explanation of an automated decision ought to suffice for assessing its credentials."

In [8] Authors reflect on several criticisms that need to be addressed.

- *Disagreements on the scope of XAI.* As for the causes of this disagreement, authors hypothesize that both interdisciplinarity and lack of rigor may have played a role.
- *Lack of definitional cohesion, precision, and adaptation.* The title defines the criticism.
- *Misleading motivations for XAI research.* It is usually based on three statements: 1) People do not trust black box AI methods; 2) The inability to reveal their inner workings is what causes people not to trust black box AI methods; and 3) Explanations promote trust. There is insufficient evidence supporting these motivating hypotheses argue the authors.
- *Limited and inconsistent evaluations.* Although several ways to evaluating XAI methods have been proposed, no approach has been broadly adopted.

XAI as an interdisciplinary field in a mature or premature point depending on how you look at it, with a large number of publications. But, it needs to mature further in the fundamental aspects of theoretical and practical formalization. I agree XAI must evolve towards a discipline of complete utility to the important problem it addresses. It needs to explain why, what, and what, for each study or proposal.

Finally, I must highlight a progress regarding a collective discussion made by several renowned authors in the field in the following paper [5], discussing a Manifesto XAI 2.0, with the aim to define and briefly describe the open challenges in the field to face. The Manifesto is a mechanism for shaping our shared visions about research in the field of XAI, and it is the outcome of the engagement of diverse expertise and different experiences by its authors. This was summarized in nine points of interest to analyze, which raise and address as a future plan, and converges with the analysis of weaknesses and problems discussed in this section. They nine points are: 1) *Creating Explanations for New Types of AI*, 2) *Improving*

(and Augmenting) Current XAI Methods, 3) *Clarifying the Use of Concepts in XAI*, 4) *Evaluating XAI Methods and Explanations*, 5) *Supporting the Human-Centeredness of Explanations: To create human-understandable explanations*, 6) *Supporting the Multi-Dimensionality of Explainability*, 7) *Adjusting XAI Methods and Explanations*, 8) *Mitigating the Negative Impact of XAI* and 9) *Improving the Societal Impact of XAI*.

These nine points reflect many of the criticisms emphasized previously. This collective paper highlight explanations and also audience as important elements. But of course, they are still challenges that need to be addressed.

In the next three sections we discuss some of the mentioned critiques, from 3 prisms, definition, explanations and evaluation. Obviously, many other prisms and lenses need to be studied together with global reflections that approach the correct direction of investigation and advances in XAI.

IV. ON THE XAI DEFINITION

The mentioned criticism is clear. The lack of definitional cohesion in the field of explainable XAI has led to significant challenges in focusing the definition and scope of the discipline. This ambiguity limits the development of standardized methodologies and metrics, making it difficult for researchers and practitioners to evaluate and compare different XAI approaches. Consequently, the absence of a unified definition can result in fragmented efforts and hinder the progress towards achieving truly interpretable and trustworthy AI systems.

In [12] is analyzed XAI from the terms used along the literature: transparency, intelligibility, interpretability and explainability. Authors use the dictionary definitions to get a departure point.

- The word "transparent" refers to something that is "easily seen through, recognized, understood, detected; manifest, evident, obvious, clear" (Oxford English Dictionary).
- An "intelligible" system should "capable of being understood; comprehensible" (Oxford English Dictionary).
- The word "interpret" definition is "to expound the meaning of (something abstruse or mysterious); to render (words, writings, an author, etc.) clear or explicit; to elucidate; to explain" (Oxford English Dictionary).
- For the word "explain", the following definitions are used: "to provide an explanation for something to make plain or intelligible" (Oxford English Dictionary), "to make something clear or easy to understand by describing or giving information about it" (Cambridge Dictionary).

It is continuously repeated a word or idea, "understand" or "easily understood". We already have a convergent term, but we have to ask ourselves another question. Does "understand" mean the same to the designer of the AI model, to the person who uses it, or to the person who is the recipient of its usage? Let's think about the medical field, the designer, the programmer, the owner company, the doctor, the patient or society in general. It is certain that their vision of understanding an AI system is different.

In [13], it was placed audience (see Figure 1) as a key aspect to be considered when explaining an AI model. It was

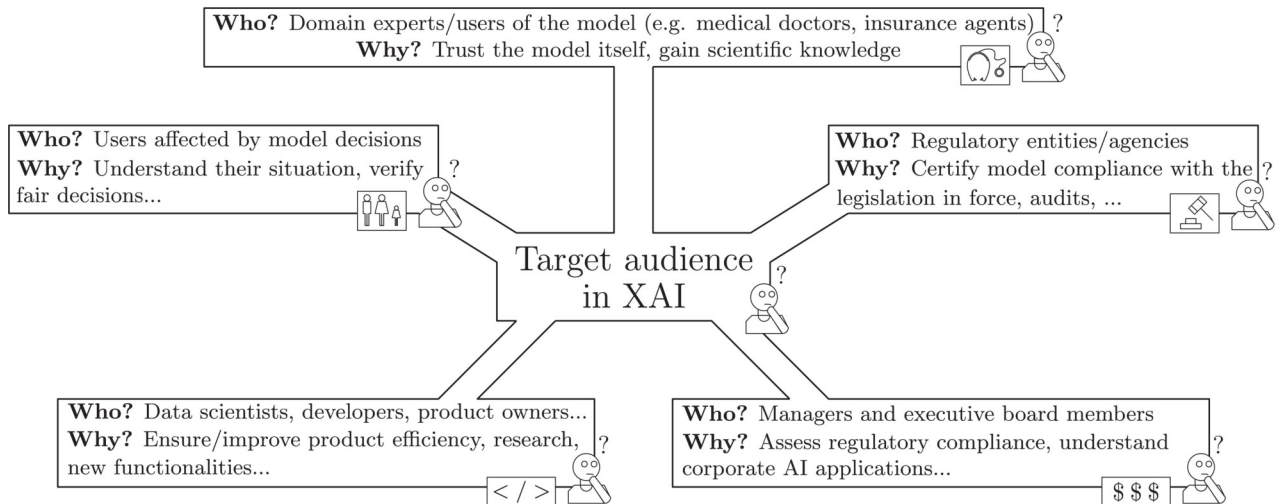


Fig. 1. Diagram showing different audience profiles (From Arrieta et al., 2220)

also elaborated on the diverse purposes sought when using XAI techniques, from trustworthiness to regulatory compliance, which round up the claimed importance of purpose and targeted audience (experts, users, developers, regulatory entities, and managers) in model explainability. In [14] was established a stakeholder interest map. It includes six levels of audience: Developer, Designer, Owner, User, Regulator, Society. This raises an interesting discussion. How would we answer the following four questions?

- Are they equal for “understanding”?
- What is their “understood” requirement?
- What is an “explanation” for them?
- What stakeholder is observing?

Under these considerations, the following definition considers both discussed elements.

Definition. [13] *Given an audience, an explainable AI is one that produces details or reasons to make its functioning clear or easy to understand.*

Regarding the lack of definitional cohesion, I believe that the definition provided in [13] serves as a solid convergence point for the topic. This definition encompasses two essential aspects, understanding and audience.

In [15], [16] have been introduced studies along which explainability approaches aim to satisfy stakeholders’ desiderata and roles from stakeholders’ desiderments. Recent studies highlight the importance of stakeholders in different areas, and in many cases involving different stakeholders, such as, autonomous systems [16], medicine [17], [18], [19] and education [20] among others.

To finish with the definition and without going into the XAI taxonomy in depth, we must distinguish between two kind of models, interpretable models versus black box AI ones [21], [13]. Models that are interpretable per se, that introduce comprehensibility on the knowledge and the inference action, for example rule base systems or decision trees with few

variables (local rule-based explainers produce logical rules which are close to human reasoning and make them suitable for non-experts). These are as opposed to black boxes, as boosting or neural networks or among others, whose difficulty of explanation increases with the neural networks number of layers. Black box AI models require post-hoc analysis. In [13], it was introduced a complete taxonomy on the post-hoc approaches with a conceptual diagram showing the different post-hoc techniques available for a machine learning (ML) model. This is an important aspect to consider in the post-hoc analysis of the black box ML models that needs connected with stakeholders’ desiderments and needs.

V. FROM DATA TYPE EXPLANATIONS TO LOCAL LINEAR EXPLANATIONS, CONCEPT-BASED EXPLANATIONS AND PROTOTYPE-BASED ONES

In the context of XAI, explanations play a pivotal role in bridging the gap between complex ML models and human understanding. By providing clear and interpretable insights into how AI systems make decisions, explanations enhance transparency, build trust, and facilitate accountability. Addressing the importance of explanations, we delve into the discussed criticisms, and we must explore various techniques and methodologies that aim to make AI models more comprehensible. We move from a local linear explanations, the most popular approach, and beyond measuring features contribution, to the general idea of explanations based on the data type, the concept-based explanations, and the use of prototypes as potential element for explain decisions in complex problems/models.

We like to mention two papers. The paper [6], it includes a deep description of some XAI techniques: SHAP: SHapley Additive exPlanations, DiCE: Diverse Counterfactual Explanations, Transformers Interpret (TI) (for language models), Grad-CAM (image classification), Layer-wise Relevance (explain image), Logic Tensor (Neural-Symbolic AI), and TS4NLE

(for natural language explanations). The paper [22], it reports extensive examples of the various explanations for each data type, highlighting similarities and discrepancies of returned explanations through. It includes a website with a software repository, called *XAI Live Survey*⁷, that authors maintain to keep pace with newly emergent methods.

Local linear explanations are among the most widely used methods in XAI. This approach involves approximating the behavior of a complex black-box model in the vicinity of a specific instance by using a simpler, more interpretable model, such as linear regression. Two well-known methods for generating local linear explanations are LIME (Local Interpretable Model-Agnostic Explanations) and SHAP (SHapley Additive exPlanations). These methods help in understanding how individual features contribute to a model's prediction for a specific instance, making it easier to validate the model's decisions. LEAF framework was proposed for the evaluation and comparison of local linear explanations, with four different metrics to evaluate different desirable qualitative aspects of explanations. In [23] authors focus on the proposal of the REVEL framework (Robust Evaluation VECTORIZED Local-linear-explanation), whose main contribution is to offer a consistent and theoretically robust analysis of the black-box generated explanations, as well as being useful at a practical level for the evaluation of explanations. REVEL takes advantage of the existing state of the art and develops a series of theoretical improvements on the generation and evaluation methods. It redefines and proposes different quantitative measures to robustly assess different qualitative aspects of the explanations.

On the other hand, a categorization based on the data type and explanation type is fundamental to structure the area and to follow the advances. In [22], authors provide an explanation-based taxonomy with a comprehensive ontology of the explanations returned, taking into account the most popular data formats and associated approaches (see section 3, Figure 1, page 1724): tabular data (Feature Importance (FI) and Rule-Based (RB)), image (Saliency Maps (SM) and Concept Attribution (CA)), text (Sentence Highlighting (SH) and Attention Based (AB)), time series (Series Highlighting Attention Based) and graphs (Node Highlighting and Edge Highlighting). It also includes two transversal approaches that we will discuss later, prototypes (the user is provided with a series of examples that characterize a class of the black box) and counterfactuals (the user is provided with a series of examples similar to the input query but with different class prediction). This overview presents an exercise carried out to address a first study on this subject. They also report the most popular Python toolkits, AIX360 [24].

As a different explanation approach, concept-based explanations offer a compelling alternative by providing a more holistic view of the model's inner workings. Concept-based explanations better resemble the way humans reason and provide more intuitive and human-understandable insights by

linking model decisions to high-level concepts. This approach helps users relate AI decisions to familiar ideas or categories, making the explanations more accessible [25]. Therefore, it has emerged as a powerful new XAI paradigm, providing model explanations in terms of human-understandable units, rather than individual features, pixels, or characters. This approach enhances the explainability by aligning explanations with concepts that are meaningful to humans. By leveraging concept-based explanations, stakeholders can gain deeper insights into the AI-assisted decision-making, making it easier to identify and address potential biases, errors, and ethical concerns. Furthermore, concept-based explanations facilitate more effective communication between AI developers and end-users, fostering trust and collaboration. As a result, concept-based explanations play a crucial role in advancing the transparency, accountability, and overall trustworthiness of AI systems in an open world. An overview on concept learning is described in [26].

The Prototype-based XAI techniques are an underutilized approaches that can provide inherently interpretable ML alternatives. Prototype selection for nearest neighbor classification has a long history in the field of ML, being highlighted as an essential tool to drive improvements in nearest neighbor techniques [27]. Prototypes must play a crucial role in the landscape of XAI, as it is discussed in [13], [28] among others, serving as a bridge between traditional explanations and concept-based explanations. Prototypes fit into this spectrum by offering concrete examples that represent typical instances of a particular class or concept. They help users understand what the model considers as a "typical" example, thereby making the model's behavior more explainable. Human reasoning is often prototype-based, using representative examples as a basis for categorization and AI-assisted decision-making. For instance, in image classification, a prototype might be a representative image that the model associates with a specific label. This can be particularly useful in identifying and understanding the characteristics that the model uses to make its decisions. By providing tangible examples, prototypes enhance the explainability of both traditional and concept-based explanations, making them a valuable tool in the XAI toolkit. We have also mentioned an approaches associated to prototype, but with with opposite use. the counterfactual explanations [29]. With a counterfactual explanation the user is provided with a series of examples similar to the input query but with different class prediction. In [13] was introduced the idea of counterfactual fairness, it tries to interpret the causes of bias using.

To finish with a challenge, XAI can be integrated as a technical objective for designers, as suggested in [30], to enhance its utility by aligning it with its intended purpose. For instance, a technical objective could be to leverage explanations to improve AI safety, as proposed in [8] with the concept of RED XAI. For example, XAI can be valuable in addressing the out-of-distribution detection problem [31], [32].

Addressing the criticisms, explanations are essential for building trust and transparency in AI systems. Concept-based

⁷<https://kdd-lab.github.io/XAISurvey/>

explanations, for instance, bridge the gap between complex models and human understanding by aligning model behavior with human-recognizable concepts. Prototype-based XAI techniques enhance interpretability by providing concrete instances that illustrate how the model operates. A comprehensive theory of explanations should encompass data, model, and post-hoc explainability, as discussed in [4].

While there may be insufficient evidence to universally support this claim, explanations play a crucial role in promoting trust in specific contexts. For instance, in high-risk scenarios like healthcare or finance, understanding the rationale behind AI decisions is essential for users to trust and accept those decisions. Additionally, explanations can help users learn from AI systems, thereby improving their own decision-making processes. Therefore, let's adopt the title of the paper by Hen et al. [33] as an aphorism and fundamental goal: "*Understanding the role of human intuition on reliance in human-AI decision-making with explanations*".

VI. ON THE MATURITY LEVEL OF EXPLAINABILITY

Assessing the maturity level of explainability techniques in XAI involves evaluating various aspects to ensure they are effective, reliable, and useful. But we have to ask ourselves, from where and how?

The maturity level of explainability in AI can be assessed through several dimensions, aligned with the AI regulation debate. I propose a triple dimension, practicality, governance and auditability. In the following are shortly described the key points:

- 1) **Practicality:** Explainability in AI is becoming more practical as tools and techniques are developed to make AI systems more transparent. This includes the creation of interpretable models and the use of surrogate models to explain complex AI systems in real problems. There are a lot of practical studies, healthcare [34], finance [35], among many others applied areas, but it is necessary a methodology for a wide practical use.
- 2) **Governance:** AI governance refers to the frameworks, processes, rules, and standards that ensure AI systems are safe, ethical, and aligned with societal values. It is crucial for several reasons, ethical development, compliance and innovation, among others. From a governance perspective, frameworks need being established to ensure that AI systems are explainable. This includes guidelines for AI governance, for ethical AI and the development of metrics to assess the explainability of AI systems [36].
- 3) **Auditability:** AI auditability refers to the ability to assess and verify AI systems' algorithms, models, data, and design processes. Explainability is also crucial for the auditability of AI systems. Being able to explain AI decisions allows for better oversight and accountability, which is essential for building trust in AI technologies [37].

There is still room for improvement in development methodologies from the above-mentioned perspectives.

VII. CONCLUDING REMARKS

Focused on the aforementioned criticisms and troubles, and from the perspective of a great theoretical development and not practical, I recognize that there is still a long way to go to enhance the usefulness of XAI. Many of these perspectives have been highlighted in this brief discussion. I have focused the attention on the XAI definition based on the audience, a fundamental element to advance toward useful XAI development.

I do not want to conclude without mentioning that some topics have been left unaddressed or barely explored. For instance, the impact of generative AI from the dual perspective of explainability and the use of large language models to enhance explainability. Additionally, the risks of overconfidence in explanations, which can increase decision-makers' tendency to rely on AI predictions even when the AI system is wrong, have not been addressed. Nor has there been a discussion on how XAI itself can be useful in guiding XAI-based model improvement, or the impact of XAI on various trustworthy AI requirements. Furthermore, a more in-depth examination of metrics is needed, including their pros and cons, and how they can advance practicality, governance, and auditability. Ensuring that AI systems are effective, reliable, and useful remains paramount.

REFERENCES

- [1] D. Castelvocchi, "Can we open the black box of AI? (News Feature)," *Nature*, vol. 538, pp. 20–23, 2016.
- [2] C. Panigutti, R. Hamon, I. Hupont, D. Fernandez Llorca, D. Fano Yela, H. Junklewitz, S. Scalzo, G. Mazzini, I. Sanchez, J. Soler Garrido et al., "The role of explainable AI in the context of the AI Act," in *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*, 2023, pp. 1139–1150.
- [3] L. Nannini, J. Alonso-Moral, A. Catala, M. Lama, and S. Barro, "Operationalizing Explainable AI in the EU Regulatory Ecosystem," *IEEE Intelligent Systems*, 2024.
- [4] S. Ali, T. Abuhmed, S. El-Sappagh, K. Muhammad, J. M. Alonso-Moral, R. Confalonieri, R. Guidotti, J. Del Ser, N. Díaz-Rodríguez, and F. Herrera, "Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence," *Information fusion*, vol. 99, p. 101805, 2023.
- [5] L. Longo, M. Brcic, F. Cabitza, J. Choi, R. Confalonieri, J. Del Ser, R. Guidotti, Y. Hayashi, F. Herrera, A. Holzinger et al., "Explainable Artificial Intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions," *Information Fusion*, vol. 106, p. 102301, 2024.
- [6] T. Freiesleben and G. König, "Dear XAI community, we need to talk! Fundamental misconceptions in current XAI research," in *World Conference on Explainable Artificial Intelligence*. Springer, 2023, pp. 48–65.
- [7] J. Zerilli, "Explaining machine learning decisions," *Philosophy of Science*, vol. 89, no. 1, pp. 1–19, 2022.
- [8] R. O. Weber, A. J. Johs, P. Goel, and J. M. Marques-Silva, "XAI is in trouble," *AI Magazine*, 2024.
- [9] N. Díaz-Rodríguez, J. Del Ser, M. Coeckelbergh, M. L. de Prado, E. Herrera-Viedma, and F. Herrera, "Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation," *Information Fusion*, vol. 99, p. 101896, 2023.
- [10] X. Huang and J. Marques-Silva, "On the failings of shapley values for explainability," *International Journal of Approximate Reasoning*, p. 109112, 2024.
- [11] J. Marques-Silva and X. Huang, "Explainability is not a game," *Communications of the ACM*, vol. 67, no. 7, pp. 66–75, 2024.

- [12] M. A. Clinciu and H. F. Hastie, "A survey of explainable AI terminology," in *1st Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence 2019*. Association for Computational Linguistics, 2019, pp. 8–13.
- [13] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins *et al.*, "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information fusion*, vol. 58, pp. 82–115, 2020.
- [14] K. Haresamudram, S. Larsson, and F. Heintz, "Three levels of AI transparency," *Computer*, vol. 56, no. 2, pp. 93–100, 2023.
- [15] M. Langer, D. Oster, T. Speith, H. Hermanns, L. Kästner, E. Schmidt, A. Sesing, and K. Baum, "What do we want from Explainable Artificial Intelligence (XAI)?—A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research," *Artificial Intelligence*, vol. 296, p. 103473, 2021.
- [16] R. R. Hoffman, S. T. Mueller, G. Klein, M. Jalaeian, and C. Tate, "Explainable ai: roles and stakeholders, desirements and challenges," *Frontiers in Computer Science*, vol. 5, p. 1117848, 2023.
- [17] H. V. Subramanian, C. Canfield, and D. B. Shank, "Designing explainable ai to improve human-ai team performance: a medical stakeholder-driven scoping review," *Artificial Intelligence in Medicine*, p. 102780, 2024.
- [18] M. Kim, S. Kim, J. Kim, T.-J. Song, and Y. Kim, "Do stakeholder needs differ?—designing stakeholder-tailored explainable artificial intelligence (xai) interfaces," *International Journal of Human-Computer Studies*, vol. 181, p. 103160, 2024.
- [19] M. Bergquist, B. Rolandsson, E. Gryska, M. Laesser, N. Hoefling, R. Heckemann, J. F. Schneiderman, and I. M. Björkman-Burtscher, "Trust and stakeholder perspectives on the implementation of ai tools in clinical radiology," *European Radiology*, vol. 34, no. 1, pp. 338–347, 2024.
- [20] A. J. Karran, P. Charland, J. Martineau, A. O. de Guinea, A. Lesage, S. Senecal, and P.-M. Leger, "Multi-stakeholder perspective on responsible artificial intelligence and acceptability in education," *arXiv preprint arXiv:2402.15027*, 2024.
- [21] M. Atzmueller, J. Fürnkranz, T. Kliegr, and U. Schmid, "Explainable and interpretable machine learning and data mining," *Data Mining and Knowledge Discovery*, pp. 1–25, 2024.
- [22] F. Bodria, F. Giannotti, R. Guidotti, F. Naretto, D. Pedreschi, and S. Rinzivillo, "Benchmarking and survey of explanation methods for black box models," *Data Mining and Knowledge Discovery*, vol. 37, no. 5, pp. 1719–1778, 2023.
- [23] I. Sevillano-García, J. Luengo, and F. Herrera, "Revel framework to measure local linear explanations for black-box models: Deep learning image classification case study," *International Journal of Intelligent Systems*, vol. 2023, no. 1, p. 8068569, 2023.
- [24] V. Arya, R. K. Bellamy, P.-Y. Chen, A. Dhurandhar, M. Hind, S. C. Hoffman, S. Houde, Q. V. Liao, R. Luss, A. Mojsilović *et al.*, "One explanation does not fit all: A toolkit and taxonomy of AI explainability techniques," *arXiv preprint arXiv:1909.03012*, 2019.
- [25] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas *et al.*, "Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav)," in *International conference on machine learning*. PMLR, 2018, pp. 2668–2677.
- [26] E. Poeta, G. Ciravegna, E. Pastor, T. Cerquitelli, and E. Baralis, "Concept-based explainable artificial intelligence: A survey," *arXiv preprint arXiv:2312.12936*, 2023.
- [27] S. Garcia, J. Derrac, J. Cano, and F. Herrera, "Prototype selection for nearest neighbor classification: Taxonomy and empirical study," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 3, pp. 417–435, 2012.
- [28] A. Narayanan and K. Bergen, "Prototype-Based Methods in Explainable AI and Emerging Opportunities in the Geosciences," in *ICML 2024 AI for Science Workshop*, 2024.
- [29] R. Guidotti, "Counterfactual explanations and how to find them: literature review and benchmarking," *Data Mining and Knowledge Discovery*, pp. 1–55, 2022.
- [30] V. Chen, J. Li, J. S. Kim, G. Plumb, and A. Talwalkar, "Interpretable machine learning: Moving from mythos to diagnostics. queue 19, 6 (jan 2022), 28–56," 2022.
- [31] H. Liu, V. Lai, and C. Tan, "Understanding the effect of out-of-distribution examples and interactive explanations on human-ai decision making," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW2, pp. 1–45, 2021.
- [32] J. Choi, J. Raghuram, R. Feng, J. Chen, S. Jha, and A. Prakash, "Concept-based explanations for out-of-distribution detectors," in *International Conference on Machine Learning*. PMLR, 2023, pp. 5817–5837.
- [33] V. Chen, Q. V. Liao, J. Wortman Vaughan, and G. Bansal, "Understanding the role of human intuition on reliance in human-ai decision-making with explanations," *Proceedings of the ACM on Human-computer Interaction*, vol. 7, no. CSCW2, pp. 1–32, 2023.
- [34] M. Frasca, D. La Torre, G. Pravettoni, and I. Cutica, "Explainable and interpretable artificial intelligence in medicine: a systematic bibliometric review," *Discover Artificial Intelligence*, vol. 4, no. 1, p. 15, 2024.
- [35] P. Weber, K. V. Carl, and O. Hinz, "Applications of explainable artificial intelligence in finance—a systematic review of finance, information systems, and computer science literature," *Management Review Quarterly*, vol. 74, no. 2, pp. 867–907, 2024.
- [36] M. A. Camilleri, "Artificial intelligence governance: Ethical considerations and implications for social responsibility," *Expert systems*, vol. 41, no. 7, p. e13406, 2024.
- [37] L. Waltersdorfer, F. J. Ekaputra, T. Miksa, and M. Sabou, "AuditMAI: Towards An Infrastructure for Continuous AI Auditing," *arXiv preprint arXiv:2406.14243*, 2024.

How CIs can Tackle Future Pandemics A Multi-Domain Approach to Improve CI Resilience

Stefan Schauer*, Manuel Egger*, Max Kesselbacher-Pirker*, Isti Rodiah†, Olga Horvadovska†, Berit Lange†, Norman FRM Fauster‡, Hannes Zenz‡, and Christian Kimmich‡

* AIT Austrian Institute of Technology
Giefinggasse 4, 1210 Vienna, Austria

Email: {stefan.schauer | manuel.egger | max.kesselbacher-pirker}@ait.ac.at

† Helmholtz Centre for Infection Research

Inhoffenstraße 7, 38124 Braunschweig, Germany

Email: {isti.rodiah | olga.hovardovska | berit.lange}@helmholtz-hzi.de

‡ Institute for Advanced Studies

Josefstädter Straße 39, 1080 Vienna, Austria

Email: {norman.fauster | hannes.zenz | christian.kimmich}@ihs.ac.at

Abstract—In this paper, we present an integrated approach that combines data and information sources from different domains to better capture the potential effects of a pandemic and to improve preparedness of critical infrastructures and decision makers in the future. This approach not only takes epidemiological data on a pathogen into account but also allows to simulate the cascading effects of the pandemic itself as well as the mitigation measures might have on the operation of CIs from various domains and, consequently, on the well-being of the society. Additionally, these effects can influence the operational capacity and economic well-being of CIs. Hence, the approach also projects the possible economic effects, i.e., monetary costs, a future pandemic might impose on society, including wide-ranging counter measures such as school closures or lock-downs.

I. INTRODUCTION

UNLIKE any other event, the COVID-19 pandemic has shown the complex and highly sensitive interrelations among the society, the critical infrastructures (CIs) and the decision makers on a national and supra-national scale. The pandemic not only had a huge impact on people's health as well as on public health but also on the functioning of critical services and thus on the social well-being of a large part of the European population. Additionally, the measures taken to mitigate the pandemic, ranging from social distancing to complete lock-downs, came with huge challenges and high cost for CI operators and national governments. From pandemic plans existing before COVID-19, it becomes evident that such wide-ranging effects and large-scale impacts were not foreseen by decision makers. However, similar pandemic events will become more likely in Europe in the future, particularly when considering climate change [4]. Hence, it is of utmost importance to prepare decision makers, CI operators and the society as a whole for future pandemics to increase their individual and combined resilience.

As part of this preparation activities, the SUNRISE project [29], funded by the European Union in the course of the Horizon Europe Programme, aims at developing a comprehensive strategy for CI operators as well as national and regional authorities to improve their robustness and resilience against future pandemic scenarios. To achieve that, the project focuses on the integration of several simulation approaches and tools from different domains such as health and epidemiology, regional and national economics as well as general aspects of CI protection. This integrated approach provides a holistic overview on the current pandemic situation to decision makers on a regional and national level as well as to CI authorities and operators.

In this paper, we will give a first insight into this integrated approach and present its overall methodological setup, i.e., the SUNRISE Process. This process represents a step-by-step guideline for CI operators and decision makers from regional and national authorities on how to prepare for and tackle an upcoming pandemic. We will show how the process utilizes data sources from different domains (i.e., epidemiological data on the pathogen, structural data on the CIs relevant for or affected by the pandemic as well as data on economic effects of the pandemic) and combines them into a decision making framework. As an example, we will also describe three simulation tools, one for describing the spreading of a virus during a pandemic, one for indicating the cascading effects of the pandemic across various industry sectors and domains and one for capturing the short- and mid-term economic effects on individual sectors.

The rest of the paper is structured as follows: in the next section, we provide a short overview on related work on pandemic preparedness and approaches to increase the protection and resilience of CIs during a pandemic. In Section III, we describe on a high level the SUNRISE approach and the

different process steps that CI operators can use to increase their resilience. As these steps are driven by ICT tools, we show some examples of simulation tools in Section IV, which can be used – and combined – to obtain a better overview on the effects of a pandemic on the CIs, the economy, and the society as a whole. Finally, Section V concludes the paper and provides an outlook on next steps in the project.

II. RELATED WORK

Critical infrastructures (CIs) are interdependent in many ways. First and foremost, CIs provide goods and services that are used by other CIs, e.g., a hospital needs electricity and water for operation, but also depends on the transportation system for staff and medication. In recent years, digitalization induced further dependencies, e.g., by electronic control systems for physical processes. Due to these interdependencies, CIs cannot be treated as isolated entities. In particular, any risk analysis carried out by an individual CI needs to take the interdependencies with other CIs into account, since those relations affect the operation of the CI itself. Furthermore, when looking at the complex network of CIs within a region or on a national scale, it is important to consider this entire network of CIs because the interdependencies affect not only risk level (i.e., the impact of particular threats on all CIs) but also the resilience (i.e., how fast all CIs can recover from an incident) of the entire network. Hence, many approaches to analyse these cascading effects such as the Cross Impact Analysis (CIA) [31], the Hierarchical Holographic Model (HHM) [9], Input-output-Interoperability Model (IIM) [23] or approaches using Interdependent Markov Chains (IDMCs) [32] have been developed. More specific models are focused on the coupling of two different domains, e.g., a power network and an ICT network by a co-simulation approach [6].

Apart from analysing cascading effects, a major challenge is to feed their consequences into the CIs' risk and resilience management. Although resilience concepts have been discussed for power distribution [18], railway transportation [5] and water distribution systems [28], amongst other sectors, they do not sufficiently address cascading effects. Recently, a combined risk and resilience management process has been proposed using a cross-domain simulation approach to integrate the consequences of cascading effects [25]. However, this process is highly generic whereas more precise guidelines for CI operators and authorities are required, tailored to the complex and wide-ranging effects of a pandemic.

A pandemic is usually assessed in terms of its effects on individual and public health, mostly by analyzing disease burden and excess mortality. For example, during the COVID-19 pandemic epidemiological and modeling studies were able to assess early on the direct disease burden by providing data or estimates on potential unreported cases [33], transmission parameters [2] case fatality and number of deaths and expected population mortality [21] and potential health care burden such as bed capacity pressures in ICU and hospital wards [11]. Nevertheless, the actual disease burden and expected excess mortality during the COVID-19 pandemic depends

on more than these estimates. In a conceptual model, other dependencies would include the direct COVID-19 burden (as described above), the indirect COVID-19 burden resulting from pressures on health system capacity other disease burden due to economic effects of the pandemic/the response to the pandemic as well as other disease burden due to social distancing as a result of the pandemic or the response to the pandemic [7]. Hence, the effects of a pandemic cannot only be assessed from an epidemiological perspective but need to include a broad variety of domains and thus need to cover a multiple impact categories. Thus, even just for the critical infrastructures of health care effects and impacts from domains beyond simply pandemic spread have to be considered. Besides healthcare systems, CIs in general are at high risk of destabilisation by both pandemic spread and anti-pandemic measures [26].

Regarding the consideration of interdependencies in (macro)economic analyses, a wide range of different models has been used for the economic impact assessment of different types of disasters. Most commonly, economic measures are applied to quantify the relations between CIs and other sectors. Inoperability input-output models (cf. [27] allow for a reduction of the operational level if intermediary inputs are not available. Other approaches may for example take Computable General Equilibrium (CGE) models or network approaches [30] into account, each with different advantages and drawbacks (cf. [17], [8], [16]). Alternatives might look at CIs as individual agents, interconnected by a set of relations. This perspective allows applying agent-based macroeconomic models developed for assessing the impact of natural disasters [3] or pandemics [20] to CIs, too.

III. SUNRISE PROCESS

The SUNRISE approach is described in the form of an iterative, step-by-step guideline, i.e., the SUNRISE Process, describing the individual activities that can support CI operators in increasing the resilience of their respective infrastructures. The SUNRISE Process is based on existing principles and standards such as the PDCA (Plan, Do, Check, Act) Cycle and the International Organization for Standardization (ISO) 31000 [12] standard for risk management. In this way, the SUNRISE Process implicitly builds on concepts, structures and mechanisms that are already existing within CIs as well as regional and national governmental organisations. Accordingly, the SUNRISE Process consists of the five major building blocks “Establishing the Context”, “Assessing the Pandemic”, “Analysing the Consequences”, “Evaluating the Measures” and “Evaluating the Resilience” (see also Fig. 1).

The first block, *Establishing the Context*, sets the scene for the SUNRISE Process and the core aspects for implementing the process are defined. First, this includes the identification of the stakeholders, i.e., the people that are interested in and benefit from the process in general and from its results, in particular. Among them are also the relevant Pandemic-Specific Critical Entities (PSCEs), which are services, infrastructures or people that are mostly affected by the different consequences of a pandemic. As the relation among the PSCEs are of




Fig. 1. Illustration of the SUNRISE Process

high interest in the SUNRISE Process, these relations and interdependencies are specifically captured in the next part of the context establishment. These interrelations will also be the basis for analysing the impacts and cascading effects a pandemic can have across different industry sectors and domains of social life. As final steps of this first block, the available data sources and general requirements are identified.

The second block is dedicated to gather information about the main threat, i.e., the pandemic that the CI or the regional or national government is facing. Therefore, the pathogen must be identified in the beginning, which can be done, for example, by using national or international surveillance and monitoring systems. Once it is clear which pathogen is causing the pandemic, more information about its characteristics is required, such as transmissibility, exposure, seriousness of disease, Case Fatality Rate (CFR) and others. These characteristics are essential to obtain a better estimation on the spreading of the pathogen and to decide on possible measure to protect from infection or reduce the spreading. The second block concludes with the definition of potential scenarios that the organisation implementing the process could be facing.

When the scenarios are described in detail, the consequences of the pandemic are analysed in the third block of the SUNRISE Process. Since one major objective of the process is to

capture multi-criteria impacts, the consequences are analysed according to four domains, i.e., the effects on the population, on PSCEs, on the economy in general and on the society as a whole. In this way, the process makes sure that the impacts of a pandemic together with the available countermeasures are not only analysed according to the effects on individual and public health but also effects on vital services, economic processes and the societal well-being is captured as well. This multi-criteria approach is of particular importance for governmental organisations on a regional and national level to make sure that they obtain a holistic overview on the impacts of a pandemic and can also identify the best countermeasures not only according to one indicator but to several indicators.

After getting an estimation on the consequences, the fourth block of the SUNRISE Process deals with the identification and evaluation of possible measures to prevent, protect against or mitigate the pandemic. As the SUNRISE Process is focusing on a multi-criteria analysis, also the countermeasures are gathered from different domains: non-pharmaceutical interventions (NPIs) on a personal, environmental and populational level, economic measures and legal measures. The NPIs focus mainly on the health of individual people, i.e., how to protect someone from getting infected with the pathogen or curing their illness, sa well as on reducing the spreading of the pathogen in the general society. Hence, some NPIs such as school closures or lockdowns potentially have huge effects onto the society and implications for the daily life, which need to be taken into account. Since most of the NPIs come with a high cost that cannot be covered by individual organisations, the economic measures describe actions how a state can help in this context, e.g., by providing funds or financial support. All of the measures taken also need to be set within a legal framework as laws and directives are still valid in the course of a pandemic.

The final block of the SUNRISE Process now covers the estimation of the risk level and the resilience level of the services, infrastructures and population in the focus of the analysis. Therefore, the data coming from the consequence analysis is gathered and compiled into one abstract level representing the risk for a given scenario, e.g., a value between 1 and 5 on a semi-quantitative risk scale. The same is done for the resilience level; here, the resilience of individual services and infrastructures is compiled into a resilience level for an entire region or nation. As a second step in this block, the various countermeasures from the previous block are taken into account and a “what-if” analysis is carried out. This analysis assumes that one or several of the measures are implemented and re-calculates the consequences with these measures in place. This will result in a new risk and resilience level, giving the decision makers an estimation, on which set of measures will be most effective according to the criteria from the different domains.

IV. SUNRISE SIMULATION TOOLS

The individual steps of the SUNRISE Process can be implemented in different ways, either by literature review and

research of existing data sources (e.g., for establishing the context or characterizing the pathogen), by bringing together experts from the various fields and sharing their knowledge in a workshop setting (e.g., for scenario description) or by the application of existing tools for the respective domains. In particular when it comes to analyzing the potential consequences of a pandemic (i.e., Step 3 of the SUNRISE Process, as described in Section III), there are several data sources and specific tools at hand that can support these tasks.

In the following, we will give three indicative examples of tools that facilitate the analysis of the spreading of a pandemic, the resulting economic impacts as well as cascading effects on CIs from various domains and have been extended and adapted in the course of the SUNRISE project. Additionally, there are more tools under development in the SUNRISE project covering other aspects of the analysis of consequences of a pandemic. In particular, four tools are implemented in the project providing specific technical solutions for mitigation activities and the support of NPIs, which are part of Step 4 of the SUNRISE Process. However, a complete description of all these tools would go beyond the scope of this paper.

A. Epidemiological Simulation

The multi-patch epidemiological model is a computational framework used to simulate the spread of infectious diseases in a spatially heterogeneous environment. Unlike simple models that assume homogeneity in population distribution, the multi-patch model acknowledges the spatial heterogeneity in populations, dividing them into multiple interconnected patches or compartments. Each patch represents a distinct geographic area or population subgroup where disease transmission can occur.

The model simulation was originally developed by Rodiah [22], [10] using Python. Parameters, initial conditions, and connectivity matrices can be specified in standard data formats such as Excel, CSV, or custom text files. Simulation results are typically saved as time-series data or visualizations, including graphs and heatmaps. Output formats may include CSV files for data analysis, image files for visualizations. The computational resources required depend on the complexity of the model and the scale of the simulation. Simulations involving many patches or detailed spatial resolution may require significant computational resources, including high-performance computing clusters or cloud-based infrastructure. Memory and processing power are essential considerations, particularly for simulations with a large number of compartments.

At the core of the multi-patch model are differential equations that describe the flow of individuals between patches and the transmission dynamics of the disease within each patch. These equations incorporate parameters such as transmission rates, recovery rates, and movement rates between patches, which are crucial in understanding how the disease spreads across different locations and populations.

One key aspect of the multi-patch model is its ability to capture the effects of spatial connectivity on disease transmission. By considering movement between patches, the model




Fig. 2. Illustration of the epidemiological model for direct transmission.

can account for the flow of infected individuals from one location to another, potentially leading to the introduction or amplification of the disease in new areas. This spatial perspective is particularly relevant for diseases with long incubation periods or those transmitted by mobile hosts, such as humans or animals.

Furthermore, the multi-patch model allows for the exploration of spatial heterogeneity in factors influencing disease transmission, such as population density, contact patterns, and environmental conditions. These variations can have significant impacts on the spread and persistence of infectious diseases, making it essential to consider spatial dynamics in epidemiological modelling and control strategies.

The model is typically set up as a system of differential equations, where each patch is represented by a set of state variables describing the population dynamics within that patch. For the SUNRISE project, each patch within the model corresponds to a Nomenclature of Territorial Units for Statistic Level 1 (NUTS1) subdivision. Within each patch, a meta-population framework is employed to account for the heterogeneity of populations across different CIs. This approach considers various demographic factors, such as age distribution and contact patterns, to capture the nuances of disease transmission within and between subpopulations.

The epidemiological dynamics within each patch are developed by adapting a deterministic Susceptible-Exposed-Infectious-Recovered (SEIR) model, including those transmitted through direct contact or vector-borne transmission. This model distinguishes between healthy (susceptible) individuals, infected but not yet infectious (exposed) individuals, and infectious patients. Moreover, depending on the nature and severity of the disease, it is possible to introduce additional compartments. In the case of severe illness, compartments for hospitalized patients and individuals in intensive care units (ICUs) can be integrated into the model. Subsequently, patients may either recover or die from the disease. Furthermore, an additional compartment is introduced to account for the indirect impact of the epidemic on critical infrastructures, represented as an absence compartment. The model structure within each patch is illustrated in Fig. 2. Therein, an individual in meta-population i is classified either as susceptible (S_i), absence (A_i), exposed (E_i), infectious (I_i), hospitalized (H_i), in intensive care (U_i), recovered (R_i), or dead (D_i).

In scenarios involving vector-borne transmission, the model incorporates additional compartments and parameters to represent the dynamics of the vector population, as well as the transmission dynamics between vectors and humans. This




Fig. 3. Illustration of the epidemiological model for vector-borne transmission.

entails introducing compartments for susceptible vectors, exposed vectors, and infectious vectors, alongside parameters governing transmission rates between vectors and humans. The model integrates the interactions between human and vector populations to simulate the spread of the disease in such transmission scenarios. Fig. 3 illustrates the model structure for vector-borne transmission, where a human individual in meta-population i is classified either as susceptible (S_{h_i}), exposed (E_{h_i}), infectious (I_{h_i}), or recovered (R_{h_i}). A vector in meta-population i is classified either as susceptible (S_{v_i}), exposed (E_{v_i}), or infectious (I_{v_i}). Transmission between human and vector represents by red dash line.

B. Cascading Effects Simulation

A Cascading Effects Simulation and Risk Analysis application (in short: CASSANDRA) has been developed in the course of the SUNRISE project, building on and extending already existing approaches of this software. The core of the simulation tool is a NodeJS-based package covering a stochastic simulation of the cascading effects an incident might have on a network of network of assets (which corresponds to a network of interrelated CIs in SUNRISE). Further, the tool has an Angular-based web front-end to support the modelling and to display the simulation results accompanied by a NestJS application providing a REST-API as an endpoint for automated access to the simulation back-end.

The first step in the analysis of cascading effects is to formally describe and model the network of CIs and to create an interdependency graph out of that. In general, this can be broken down into two main parts, i.e., identification of relevant components and identification of the dependencies between these components. The identification of the relevant components depends on the purpose of the analysis. If the focus lies on raising awareness of various existing dependencies, a high-level diagram is sufficient where each CI is represented as one node. If the big picture is known and the focus lies on a deeper understanding, it is required to model a CI in more detail, i.e., represent all its relevant (critical) components as nodes. The granularity depends on the purpose of the analysis, and in some cases also on the availability of data.

When focusing on the effects of an incident, and particularly on cascading effects within a CI network, it is necessary to understand the direction of the propagation of these effects. In the interdependency graph, this is realized by using a directed

graph. In the context of the interdependency graph, an edge $X \rightarrow Y$ means that a problem in component X may influence component Y . For example, hospital Y needs drinking water from water utility X .

In general, the dependencies can be of various kinds, and the type of dependency may influence the propagation in the sense that the probability that the problem affects other components may depend on it [15]. An alternative modelling approach classifies the nodes of the networks as physical, cyber, process, human etc. and characterizes the propagation through the node's behaviour [13]. In case of the above mentioned hospital depending on multiple products of the water provider, drinking and cooling water, this dependency is represented as one physical dependency. If this dependency is important, e.g., if the hospital is in the focus of the analysis, a more detailed representation is preferred.

The main purpose of the interdependency graph is to obtain information about the global behaviour of the CI network based on the local behaviour of CIs. These local dynamics are described through a model inside each node, describing how a threat affects this specific node. The first task is therefore to measure this effect. Due to the complexity of the modelled components (either entire CIs or their critical components), it is not feasible to use specific and detailed measures of loss for each node. Instead, a qualitative scale to characterize the state of the node is more favourable, e.g., ranging from 1 (best) to 5 (worst). Depending on the type of the node, the levels represent functionality or availability of a component.

In the context of CIs or their crucial components, data is often sparse or vague, which makes a precise and detailed description of the local dynamics almost impossible. With the choice of a qualitative state, specification of the node's dynamic boils down to describing when it changes its state, i.e., when the condition gets better or worse. Such a change is triggered by an incident, either directly or indirectly through the state of a node it depends on. Further, the reaction to a threat may depend on the circumstances, i.e., on the current state of the node. Such behaviour is best modelled through a Mealy automaton, as it changes its state upon a given input and returns an output. The reaction of a node to an input is influenced in real life by manifold factors that can hardly be captured in full detail in a practical abstract model. Therefore, it is appropriate to model a node's behaviour by adding probabilities to the state changes of the automaton model, i.e., through a probabilistic Mealy automaton [14].

Based on the local dynamics of the individual CIs within the interdependency graph, a simulation approach can be used to describe the global dynamics, i.e., the behaviour of the entire CI network upon a specific incident happening at one of the nodes. This is realized by sending notifications from one node to all its neighbours if a problem has occurred. The Mealy automaton inside each node reacts to an input α and returns an output β if it changes the state (i.e., if it is affected by the trigger). All neighbouring nodes receive this output as new input and may react accordingly. Through this transmission of messages which can be interpreted as alarms, the impact of

an incident can propagate through the entire network.

This path of events through all possible dependencies modelled in the interdependency graph describes the potential cascading effects of the threat affecting the initial node. This simulation is carried out by a tool developed by AIT [24], [1], which implements this stochastic process. Because of the probability distributions of the state transitions in the individual nodes, each simulation could lead to a different result. The overall impact of the cascading effects on the entire CI network is then measured by the resulting states of the individual CIs. Hence, the tool runs numerous iterations of the simulation to get a statistical overview on the results.

For the SUNRISE project, to provide a multi-faceted view on cascading pandemic effects, we develop and continuously refine an interdependency graph that can be integrated with both the epidemiological (cf. Section IV-A) and economic (cf. Section IV-C) simulations. To facilitate the integration of the epidemiological simulation, CIs are connected to: i) regional nodes based on their NUTS2 region, and ii) population nodes divided by pandemic-specific age groups (e.g. children, adults, elderly). These connections make it possible to setup transitions for pandemic events like a threshold of the population being admitted to hospitals or ICUs, in turn affecting the operational level of the CIs. In this case, the cascading effects simulation uses simulation data from the epidemiological simulation.

Regarding the integration of the economic simulation, CIs are connected to nodes representing NACE sectors (division of sectors based on [19]). The CIs are connected to specific sectors based on their production and demand of goods and services. These connections make it possible to transfer effects of pandemic events to national NACE sectors, affecting the operational level of economic sectors. In this case, the cascading effects simulation provides input data to the economic simulation in the form of the pandemic-related degrading of operational levels of economic sectors.

C. Economic Impact Simulation

To capture the economic impact of a pandemic, an agent-based model (ABM) has been developed in the course of the SUNRISE project, building on a sector-disaggregated macro-economic model originally created by Poledna et al. [20]. One crucial aspect in SUNRISE was the required flexibility of the analysis, since the model should be able to potentially cover a wide range of sectors and disasters. The model agents form expectations in each simulation period regarding income, demand, and growth of the Gross Domestic Product (GDP) amongst others based on an autoregressive process of order one. Thus, the agents are not equipped with rational nor model-consistent expectations.

The ABM is implemented as a Matlab simulation and was originally developed by [20] in an open-access manner. Disaster- or model-specific inputs like shocks can be considered in several ways. One possibility is importing personnel numbers, productivity losses or similar through Excel or CSV files.

A complete model computation corresponds to a Monte Carlo simulation of individual model runs. One single model run consists of the iteration through the pre-set timesteps (in quarters), computing all prices, investments, expenditures and others. Random processes are added to the expectations of economic growth and prices, imports, exports, government consumption, and shocks. For each Monte Carlo step, these random elements are newly drawn and the aggregated macro-economic and disaggregated sectoral indicators computed. For the final results, the variables are summarized over all runs to average out the effects of the random components.

The simulation approach can account for CIs and related capital stocks of sectors based on the statistical classification of economic activities in the European Community, i.e., the NACE level (short for nomenclature statistique des activités économiques dans la Communauté européenne) of the Figaro tables (i.e., Full international and global accounts for research in input-output Analysis) provided by Eurostat. The model is thus based on an input-output framework and originally calibrated for the small open economy of Austria. For the SUNRISE project, the model is being continuously refined and adapted to other national economies and pandemic scenarios. With the given model architecture, data for other European countries can be used for calibration which makes a simulation for those economies possible as well.

The economic impact simulation in SUNRISE considers the following sectors: firms, private households, the general government, banks including the central bank, and the rest of the world. Each sector consists of heterogeneous agents representing either natural persons or legal entities that interact according to predefined rules (see Figure 4). The firm sector is made up of 64 industries, each producing a perfectly substitutable good with labour, capital, and intermediate inputs from other sectors with a fixed-coefficients technology. The model is based on quarterly data and typically runs simulations for up to three years in the current version; this implies a forecasting period of 12 quarters. The model architecture is flexible and allows for several types of simultaneous shocks on a sectoral level. Examples include supply shocks (e.g., due to disruptions in the supply chain), demand shocks (e.g., travel restrictions), changes in productivity (e.g., employees absent from work due to an infection, quarantine or because they need to take care of others) or destruction of capital stock. It is further easy to change parameters to assess the implications and impacts throughout the modelled economy.

The main data source of the ABM is economic data including input output tables, national accounts, capital stocks, business demography, government statistics and population data mainly provided by Eurostat. For the simulation of various pandemic scenarios, additional inputs are required, such as information on the number of persons absent from work (coming from the epidemiological model, cf. Section IV-A, estimated reduction of sectoral output, changes regarding the service level (both coming from the cascading effects simulation, cf. Section IV-B), counter-measures like lock-downs or travel restrictions and others. As an output, the ABM provides




Fig. 4. Illustration of the model agents and their interactions

the standard indicators of a macroeconomic model on which the socio-economic impacts are assessed, e.g., the nominal and real gross value added as well as the employment in each considered sector, the GDP, the total employment as well as the unemployment rate. However, given the model architecture, it is easy to define new economic output variables and have them displayed in addition to the standard macroeconomic outputs.

V. CONCLUSION

In the face of a pandemic, it is important to support decision makers on a regional and national level as well as CI operators on the selection of the most effective counter measures. However, the COVID-19 pandemic has shown that a sole focus on epidemiological factors is not sufficient in that case but a more holistic view is required that also takes the functionality of the CIs and the socio-economic effects of the respective measures into account. The SUNRISE Process provides such a holistic view by integrating various simulation methods and by evaluating the effects of a pandemic according to multi-domain criteria.

The SUNRISE process as described here is currently given as a first draft and will be further elaborated on in the course of the project. Next steps in the project include a more detailed specification of the process, the integration of additional simulation approaches and the validation of the overall process with CI operators and regional authorities from Italy, Spain and Slovenia.

ACKNOWLEDGEMENT

This work has been funded in the course of the SUNRISE Project as part of the European Union's Horizon Europe research and innovation programme under grant agreement no. 101073821.

REFERENCES

- [1] AIT Austrian Institute of Technology. Cascading Effects Simulation and Risk Analysis (CASSANDRA) Tool Website, 2024.
- [2] Cleo Anastassopoulou, Lucia Russo, Athanasios Tsakris, and Constantinos Siettos. Data-based analysis, modelling and forecasting of the COVID-19 outbreak. *PLOS ONE*, 15(3):e0230405, March 2020. Publisher: Public Library of Science.
- [3] Gabriel Bachner, Nina Knittel, Sebastian Poledna, Stefan Hochrainer-Stigler, and Karina Reiter. Revealing indirect risks in complex socio-economic systems: A highly detailed multi-model analysis of flood events in Austria. *Risk Analysis*, 44(1):229–243, 2024. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.14144](https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.14144).

- [4] Rachel E. Baker, Ayesha S. Mahmud, Ian F. Miller, Malavika Rajeev, Fidisoa Rasambainarivo, Benjamin L. Rice, Saki Takahashi, Andrew J. Tatem, Caroline E. Wagner, Lin-Fa Wang, Amy Wesolowski, and C. Jessica E. Metcalf. Infectious disease in an era of global change. *Nature Reviews Microbiology*, 20(4):193–205, April 2022. Publisher: Nature Publishing Group.
- [5] Nikola Bešinović. Resilience in railway transport systems: a literature review and research agenda. *Transport Reviews*, 40(4):457–478, January 2020.
- [6] Benjamin Camus, Anne Blavette, Anne-Cécile Orgerie, and Jean-Baptiste Blanc-Rouchossé. Co-simulation of an electrical distribution network and its supervision communication network. In *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*, pages 1–6, January 2020. ISSN: 2331-9860.
- [7] Florian Dorn, Berit Lange, Martin Braml, David Gstrein, John L.Z. Nyirenda, Patrizio Vanella, Joachim Winter, Clemens Fuest, and Gérard Krause. The challenge of estimating the direct and indirect effects of COVID-19 interventions – Toward an integrated economic and epidemiological approach. *Economics and Human Biology*, 49:101198, April 2023.
- [8] Luca Galbusera and Georgios Giannopoulos. On input-output economic models in disaster impact assessment. *International Journal of Disaster Risk Reduction*, 30:186–198, 2018.
- [9] Yacov Y. Haimes. Hierarchical Holographic Modeling. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(9):606–617, 1981.
- [10] Manuela Harries, Veronika K. Jaeger, Isti Rodiah, Max J. Hassenstein, Julia Ortmann, Maren Dreier, Isabell von Holt, Melanie Brinkmann, Alex Dulovic, Daniela Gornyk, Olga Hovardovska, Christina Kuczewski, Marc-André Kurosinski, Maïke Schlotz, Nicole Schneiderhan-Marra, Monika Strengert, Gérard Krause, Martina Sester, Florian Klein, Astrid Petersmann, André Karch, and Berit Lange. Bridging the gap - estimation of 2022/2023 SARS-CoV-2 healthcare burden in Germany based on multidimensional data from a rapid epidemic panel. *International Journal of Infectious Diseases*, 139:50–58, 2024.
- [11] Joel Hellewell, Sam Abbott, Amy Gimma, Nikos I. Bosse, Christopher I. Jarvis, Timothy W. Russell, James D. Munday, Adam J. Kucharski, W. John Edmunds, Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group, Sebastian Funk, and Rosalind M. Eggo. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet. Global Health*, 8(4):e488–e496, April 2020.
- [12] International Organization for Standardization (ISO). ISO/IEC 31000:2018 Risk Management - Principles and Guidelines. Technical report, Geneva, Switzerland, 2018.
- [13] Sandra König, Antonios Gouglidis, Stefan Rass, Neil Adams, Paul Smith, and David Hutchison. Analyzing Disaster-Induced Cascading Effects in Hybrid Critical Infrastructures: A Practical Approach. In Jacek Rak and David Hutchison, editors, *Guide to Disaster-Resilient Communication Networks*, pages 769–789. Springer International Publishing, Cham, 2020.
- [14] Sandra König, Stefan Rass, Benjamin Rainer, and Stefan Schauer. Hybrid Dependencies Between Cyber and Physical Systems. In Kohei Arai, Rahul Bhatia, and Supriya Kapoor, editors, *Intelligent Computing, Advances in Intelligent Systems and Computing*, pages 550–565. Springer International Publishing, 2019.
- [15] Sandra König, Stefan Schauer, and Stefan Rass. A Stochastic Framework for Prediction of Malware Spreading in Heterogeneous Networks. In Billy Bob Brumley and Juha Rönning, editors, *Secure IT Systems: 21st Nordic Conference, NordSec 2016, Oulu, Finland, November 2-4, 2016. Proceedings*, pages 67–81. Springer International Publishing, Cham, 2016.
- [16] Yasuhide Okuyama. A few good models for economic analysis of disasters: can your model handle the truth? *Chapters*, pages 30–49, 2022. Publisher: Edward Elgar Publishing.
- [17] Yasuhide Okuyama and Joost R. Santos. Disaster Impact and Input-Output Analysis. *Economic Systems Research*, 26(1):1–12, January 2014. Publisher: Routledge [_eprint: https://doi.org/10.1080/09535314.2013.871505](https://doi.org/10.1080/09535314.2013.871505).
- [18] Sadeeb Simon Ottenburger, Hüseyin Kemal Çakmak, Wilfried Jakob, Andreas Blattmann, Dmytro Trybushnyi, Wolfgang Raskob, Uwe Kühnapfel, and Veit Hagenmeyer. A novel optimization method for urban resilient and fair power distribution preventing critical network states. *International Journal of Critical Infrastructure Protection*, 29:100354, June 2020.

- [19] Anton Pichler, Marco Pangallo, R. Maria del Rio-Chanona, François Lafond, and J. Doyne Farmer. Forecasting the propagation of pandemic shocks with a dynamic input-output model. *Journal of Economic Dynamics and Control*, 144:104527, November 2022.
- [20] Sebastian Poledna, Michael Gregor Miess, Cars Hommes, and Katrin Rabitsch. Economic forecasting with an agent-based model. *European Economic Review*, 151:104306, January 2023.
- [21] Dimple D. Rajgor, Meng Har Lee, Sophia Archuleta, Natasha Bagdasarian, and Swee Chye Quek. The many estimates of the COVID-19 case fatality rate. *The Lancet. Infectious Diseases*, 20(7):776–777, July 2020.
- [22] Isti Rodiah, Patrizio Vanella, Alexander Kuhlmann, Veronika K. Jaeger, Manuela Harries, Gerard Krause, Andre Karch, Wolfgang Bock, and Berit Lange. Age-specific contribution of contacts to transmission of SARS-CoV-2 in Germany. *European Journal of Epidemiology*, 38(1):39–58, January 2023.
- [23] Joost R. Santos and Yacov Y. Haimes. Modeling the demand reduction input-output (I-O) inoperability due to terrorism of interconnected infrastructures. *Risk Analysis: An Official Publication of the Society for Risk Analysis*, 24(6):1437–1451, 2004.
- [24] Stefan Schauer, Thomas Grafenauer, Sandra König, Manuel Warum, Stefan Rass, and Stefan Rass. Estimating Cascading Effects in Cyber-Physical Critical Infrastructures. In *Critical Information Infrastructures Security*, Lecture Notes in Computer Science, pages 43–56, Linköping, Sweden, 2020. Springer International Publishing.
- [25] Stefan Schauer, Martin Latzenhofer, Sandra König, and Stefan Rass. Conceptual Approach Towards a Combined Risk and Resilience Framework for Interdependent Infrastructures. In *Proceedings of the 31st European Safety and Reliability Conference (ESREL 2021)*, pages 2161–2171, Angers, France, 2021.
- [26] Cornelia Scholz, Stefan Schauer, and Martin Latzenhofer. The emergence of new critical infrastructures. Is the COVID-19 pandemic shifting our perspective on what critical infrastructures are? *International Journal of Disaster Risk Reduction*, 83:103419, December 2022.
- [27] Roberto Setola, Stefano De Porcellinis, and Marino Sforza. Critical infrastructure dependency assessment using the input–output inoperability model. *International Journal of Critical Infrastructure Protection*, 2(4):170–178, December 2009.
- [28] Sangmin Shin, Seungyub Lee, David Judi, Masood Parvania, Erfan Goharian, Timothy McPherson, and Steven Burian. A Systematic Review of Quantitative Resilience Measures for Water Infrastructure Systems. *Water*, 10(2), February 2018.
- [29] SUNRISE. HE SUNRISE - Strategies and Technologies for United and Resilient Critical Infrastructures and Vital Services in Pandemic-Stricken Europe.
- [30] Linn Svegrup, Jonas Johansson, and Henrik Hassel. Integration of Critical Infrastructure and Societal Consequence Models: Impact on Swedish Power System Mitigation Decisions. *Risk Analysis*, 39(9):1970–1996, 2019. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.13272](https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.13272).
- [31] Murray Turoff. An alternative approach to cross impact analysis. *Technological Forecasting and Social Change*, 3:309–339, 1971.
- [32] Zhifang Wang, Anna Scaglione, and Robert J. Thomas. A Markov-Transition Model for Cascading Failures in Power Grids. In *2012 45th Hawaii International Conference on System Sciences*, pages 2115–2124, Maui, HI, USA, January 2012. IEEE.
- [33] Shi Zhao, Salihu S. Musa, Qianying Lin, Jinjun Ran, Guangpu Yang, Weiming Wang, Yijun Lou, Lin Yang, Daozhou Gao, Daihai He, and Maggie H. Wang. Estimating the Unreported Number of Novel Coronavirus (2019-nCoV) Cases in China in the First Half of January 2020: A Data-Driven Modelling Analysis of the Early Outbreak. *Journal of Clinical Medicine*, 9(2):388, February 2020.

Languages for Non-developers: What, How, Where?

Invited Talk—Extended Abstract

Juha-Pekka Tolvanen
 0000-0002-6409-5972
 MetaCase,
 Jyväskylä, Finland
 jpt@metacase.com

LANGUAGES TO RAISE THE LEVEL OF ABSTRACTION

PRODUCTIVITY has improved each time programming languages have raised the level of abstraction. This trend continues today with languages that narrow the scope they address, referred to as domain-specific languages (DSLs) [1]. However, many of these DSLs are built by developers for developers and tend to focus on the solution domain rather than the problem domain. These languages typically use text as the specification language, are built on top of IDE tools used by programmers and rely on diff and merge of files for collaboration.

In this talk, we will focus on languages that are more closely aligned with the problem domain rather than the solution domain, thereby addressing the needs of domain experts. Such languages not only raise the level of abstraction beyond programming but also enable non-developers to capture and communicate their knowledge, and, together with appropriate tools, support testing, validation, and feedback. This is important as research has consistently shown that common reasons for project failures, budget overruns, and similar issues are often related to limited understanding and formulating requirements and to the lack of user involvement. By using languages that are close to the problem domain, many typical development tasks—especially those related to requirements specification, checking, and validation—can be performed by non-developers. In many cases, the specifications created by domain experts can also be used to automatically generate code, configurations, tests, deployment instructions, and more.

CHARACTERISTICS OF NON-DEVELOPER LANGUAGES

The talk is based on a review of over 200 industry cases ([2], see Figure 1) involving the creation and use of domain-specific languages with MetaEdit+ tool [3]. MetaEdit+ enables users to create and use domain-specific modeling languages and generators.

Interestingly, most of the analyzed DSLs were developed for use by non-programmers (see Figure 2), which inspired the title of this talk. We will present examples of non-developer languages, such as those used by usability experts, safety engineers, security engineers, insurance experts and instrumentation experts. These sample languages will illustrate how they differ from traditional programming languages or DSLs




Fig. 1. Examples of DSLs reviewed

created for developers and programmers: Languages for non-developers are designed to align more closely with specific domains, representing knowledge through maps, diagrams, matrices, tables, and their combinations rather than plain text alone.

We will highlight key findings from the reviewed DSLs, including who created them (Figure 3), their size relative to standardized modeling languages like UML (Figure 4), and whether languages created for domain experts are smaller or larger than those intended for use by programmers, or if the role of the language creator influences size of the language.

HOW TO CREATE LANGUAGES FOR NON-DEVELOPERS

In the second part of the talk, we will discuss how creation of non-developer languages differs from that of programming languages. While publications on domain-specific languages typically focus on their abstract syntax, defined through meta-models or grammars, we will emphasize aspects relevant to languages used by domain experts, such as the importance of concrete syntax (e.g. following guidelines like those in [4]) and the provision of support for language use, including guidance, animation, and error and warning reporting [5]—features often




Fig. 2. Primary language user (n=45)




Fig. 3. Who implements domain-specific languages (n=100)

not covered in language specifications (see e.g. languages standardized by OMG, ITU, or The Open Group). Regarding industrially used languages, we will focus on two critical aspects of using DSLs: enabling user participation during language creation and supporting the evolution of the languages, along with the co-evolution of the work made using previous versions of the language [6].

TOOL AND PROCESS SUPPORT

Computer languages also require tools, as tools can transform precise knowledge representations made with DSLs into software code and other artifacts. We will examine the history of tools used for creating DSLs, including the latest advances in this area [7] as well as the effort to create DSLs [8]. Tools are also essential for enabling collaboration, but it is often unrealistic to apply approaches used in traditional programming, such as IDE tools and diffing and merging of files, to languages used by non-developers. Domain experts and non-developers expect tools that are easier to learn and




Fig. 4. Size of languages - in terms of size of their metamodel (n=39)

use, as they do not typically use their languages on a daily basis, unlike developers who use programming languages. Domain experts also expect that collaboration, viewing, and managing changes in the artifacts will be simpler and more closely aligned with the problem domain, rather than focusing on tracking changes through character differences within files. Not only the languages but also the entire process related to using them needs to focus on the problem domain.

We conclude by envisioning the role of non-developers in language creation and identifying situations where DSLs are most suitable, as well as pointing areas where they may not be applicable.

REFERENCES

- [1] Fowler, M. 2008. "Domain-Specific Languages", Addison-Wesley
- [2] MetaCase. 2024. "DSL of the week", <https://www.facebook.com/media/set/?set=a.2102426129807641> (accessed August 2024)
- [3] MetaCase. 2023. "MetaEdit+ 5.5 User's Guides", <https://metacase.com/support/55/manuals/> (accessed August 2024)
- [4] D. Moody. 2009. The "Physics" of Notations, IEEE Transactions on Software Engineering, vol. 35, no. 6
- [5] S. Kelly, J-P. Tolvanen. 2021. "Automated Annotations in Domain-Specific Models: Analysis of 23 Cases". STAF Workshops
- [6] J-P. Tolvanen and S. Kelly. 2023. "Evaluating Tool Support for Co-Evolution of Modeling Languages, Tools and Models", ACM/IEEE MODELS Conference companion
- [7] M. Ozkaya and D. Akdur. 2021. "What do practitioners expect from the meta-modeling tools? A survey", Journal of Computer Languages, Vo 63
- [8] J-P. Tolvanen and S. Kelly. 2018. "Effort Used to Create Domain-Specific Modeling Languages". ACM/IEEE Conference on Model Driven Engineering Languages and Systems

Mixed-Methods Study of Arabic Online Review Influence on Purchase Intention (AOCR-PI)

Ahmad Alghamdi
0000-0002-2725-6968
Department of Information
Technology
Taif University
Department of Informatics
University of Sussex
Brighton, UK
Email: aa2585@sussex.ac.uk

Natalia Beloff
0000-0002-8872-7786
Department of Informatics
University of Sussex
Brighton, UK
Email: N.Beloff@sussex.ac.uk

Martin White
0000-0001-8686-2274
Department of Informatics
University of Sussex
Brighton, UK
Email: M.White@sussex.ac.uk

Abstract—Online customer reviews (OCRs) have become vital for shoppers, aiding their purchase decisions amidst the rapid growth of user-generated content. However, limited attention has been paid to studying the impact of OCRs on the purchase intentions of Arab consumers. Therefore, applying Western online review systems to other cultures without further consideration may pose challenges. This study aims to examine how various factors of OCRs affect Arab consumers' buying intentions. Employing a mixed-methods approach, quantitative data from a survey questionnaire (633 responses) and qualitative insights from interviews (15 participants) were collected and analysed sequentially. The findings reveal that review central cues (valence, comprehensiveness, readability and images) and some peripheral cues (volume and reviewer experience) significantly influence purchase intention. By contrast, reviewer identity disclosure and reputation are not deemed important by Arab book shoppers. The semi-structured interviews validated the significance of reading OCRs before purchase, offered insights into the impact of various related factors, and revealed a new factor that is shared perspectives between the reviewer and OCR receiver. The study contributes theoretical insights and provides managerial implications for ORP developers and book publishers, aiming to enhance user experience and drive sales.

Index Terms—E-Commerce; Independent online review platform; Online book reviews; Online consumer reviews; Elaboration likelihood model.

I. INTRODUCTION

ONLINE consumer reviews (OCRs) allow consumers to gain valuable insights and evaluations of products from other customers before making a purchase decision. This abundance of information empowers customers to refine their future purchase decisions based on valuable insights. In contrast to traditional product promotions that merely emphasise product advantages, the emergence of OCRs has transformed the way consumers access information. OCRs have become a trusted source of unbiased insights, offering diverse perspectives from fellow customers. Supporting this notion, a recent study found that around 90% of participants said that OCRs influence their buying decisions [1], which directly impacts sales [2]. Understanding the role of OCRs and their influence

on purchase intentions forms the central focus of our mixed-method explanatory sequential study.

An OCR is defined as “peer-generated product evaluations on company or third-party websites” [3]. OCRs can be classified as internal or external, depending on whether they are hosted by retailers (e.g. Amazon) or independent review platforms (e.g. Goodreads), respectively [4]. Although many studies focus on OCRs and consumer buying behaviour, most examine only internal OCRs, despite the significant impact of external OCRs on product sales [5]. Additionally, the effect of OCRs on buying intention varies across cultures. Neglecting users' cultural backgrounds when developing online review systems and designing websites can affect customer evaluations and perceived value [6], [7]. However, few researchers have considered the cultural background of OCR users, limiting the generalisation of current works. Therefore, future research needs to consider these factors to improve our understanding of OCRs and enhance their effectiveness in different cultural contexts.

In summary, OCRs are essential in shaping consumer choices. However, the research has primarily focused on internal OCRs, despite the greater impact of external OCRs. Moreover, the impact of OCRs on consumers' decisions varies depending on the cultural backgrounds of consumers, which may have been overlooked in the development of review systems. To overcome these limitations, this research seeks to investigate the impact of OCRs on the purchase intention of Arab consumers. Guided by the elaboration likelihood model (ELM) [8], Hall's cultural model (HCM) [9] and Hofstede's cultural dimensions framework (HCDF) [10], the research model was built and presented in our previous paper (see [11]). This paper presents the findings from data analysis and discusses the implications derived from the study.

A sequential explanatory mixed-method approach was adopted to examine the correlation between OCR factors and purchase intention and to explain the results. First, quantitative data were collected from Arab users of Goodreads.com using online questionnaires to test the research model proposed in [11]. Then, qualitative data were used to explain the

quantitative results considering Hall's and Hofstede's cultural theories. A subset of survey respondents has been recruited to participate in semi-structured interviews for collecting the qualitative data. The sample was selected from Goodreads because it is the largest online review platform (ORP) for books [12]. Therefore, this methodology provides a thorough approach to studying the correlation between OCR factors and purchase intention among Arab users of Goodreads.com.

The remainder of the article is structured as follows: In Section II, the theoretical foundation of the research model is briefly reviewed, and the research model is provided (detailed hypotheses discussion is provided in [11]). In Section III, the research methodology is presented, including details of the sample and methods of data collection, followed by the analysis procedures of the empirical data. In Section IV, the quantitative and qualitative results of this study are reported. In Section V, a detailed discussion of the quantitative and qualitative findings is provided with theoretical and practical implications. Finally, concluding remarks are presented including the research limitations and suggestions for future works in Section VI.

II. THEORETICAL FOUNDATION AND RESEARCH MODEL

ELM suggests two decision-making routes: central (entails high cognitive efforts) and peripheral (depending on simple cues) [8], [13]. This study examines online reviews as persuasion pathways for Arab consumers, considering review content (central cues) and numerical and reviewer-related factors (peripheral cues). Additionally, considering the high-context (HC) and collectivist nature of Arab culture [9], [10], the study utilises Hofstede's uncertainty avoidance (UA) and individualism (IND) dimensions to guide hypotheses and examine their influence on how Arabs use OCRs. This combined approach aims to understand the persuasive processes in action and their impact on Arab consumer purchase intentions.

Building upon these theories, our previous paper proposed the research model presented in Fig 1, discussed the associated research hypotheses and provided an extensive review of the current literature that led to the hypotheses H1 - H8 [11].

III. METHODOLOGY

A sequential explanatory mixed-method approach was adopted. Quantitative data were collected and analysed first, while qualitative interview data were collected second to explain and expand upon the results of the initial quantitative findings [14]. This study focuses on the factors of OCRs on independent ORPs that influence purchase intention. Questionnaire data was gathered from Arab users of Goodreads.com, an online reader community, and was used to test the research model. Two reasons led to the selection of Goodreads users: first, Goodreads is the largest and most popular independent platform for book reviews, and second, books are considered experience products where consumers cannot evaluate their quality before purchase. Thereby, customers often pursue the opinions of others to reduce uncertainty related to future purchases [15], especially in high UA cultures such

as Arab culture. Accordingly, purposive sampling was employed to obtain results that could be representative of a particular segment of the online population, in this study, Arab users of OCRs.

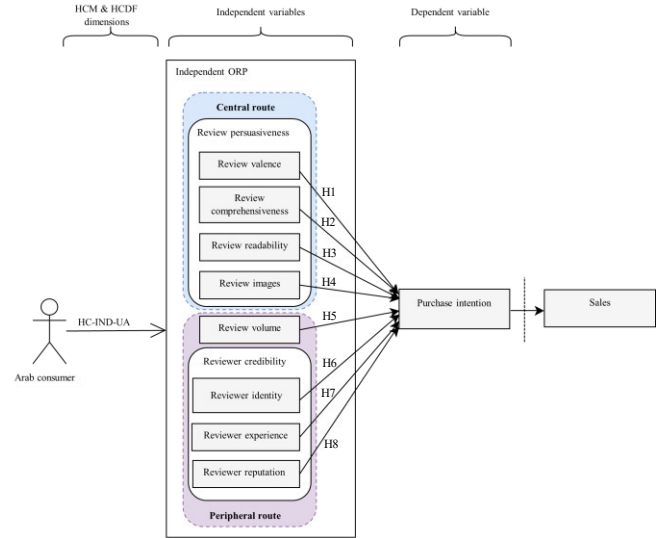


Fig 1. Research model (AOCR-PI) [11]

A. Quantitative survey design and measures (first phase)

To reach the target population, an online survey was posted in different Arabic discussion groups on Goodreads.com. Qualtrics web-based survey platform was used to design the questionnaire, which was configured to show only completed responses. The survey was initially created in English language and then provided for participants in Arabic. It was divided into two sections: the first section comprised 8 questions that covered the participant demographics and their OCR usage experience, and the second section consisted of 34 questions addressing nine research model constructs. To maintain the content validity, the questionnaire relied on established measures from past research, with modifications made to some items to fit this study. Face validity evaluation was conducted by piloting the survey with 25 participants from the target population (Arab users of Goodreads) to assess the wording and clarity of the questions. After the pilot study, some questions were clarified, and enhancements were made to the survey layout. All variables were measured using a 5-point Likert scale (from 1=strongly disagree to 5=strongly agree). The questionnaire is available from the corresponding author upon request.

For both the quantitative and qualitative studies, the Institutional Review Board approval was obtained at the University of Sussex. Furthermore, before commencing data collection, all participants in both studies provided informed written consent before completing the questionnaire and oral consent before the interview. The quantitative data was collected from May 15 to June 15, 2023.

B. Quantitative data analysis

SPSS software was used to address data screening issues and report demographic information of respondents and their OCR usage experience. Using the partial least squares (PLS) approach, the valid survey responses were entered into SmartPLS software to evaluate the survey measurements and test the research hypotheses. This method has been widely employed in information systems research [16]. PLS is a prominent method that combines factor analysis and regression analysis to simultaneously evaluate both the measurement and structural models. It is well-suited for handling complex models with many indicators [16]. Moreover, it is the most appropriate approach for predictive research [16], which is the main objective of the current study. A two-stage approach is used to execute partial least squares structural equation modelling (PLS-SEM): first, assessing the reliability and validity of the study constructs through the measurement model, and second, testing the research hypotheses by assessing the structural model [16]. Section IV describes the data analysis procedures and results.

C. Qualitative data collection and analysis (second phase)

As previously discussed, qualitative interview data were collected to provide deeper insights into the quantitative findings obtained from the questionnaires, enriching the overall understanding of the research outcomes. Accordingly, a semi-structured interview was selected, as it is suitable for exploratory and explanatory research [17]. Thus, the second part of the study was conducted using convenience sampling to interview a subset of quantitative respondents. In Arabic, the interview was piloted before conducting the final interviews to test the understandability and clarity of the questions and to estimate the average interview duration. The final semi-structured interviews were then conducted in June 2023.

During the interviews, the participants freely responded to a series of questions about online review and reviewer factors. In addition, the questions addressed the cultural characteristics of Arab OCR users (HC, IND and UA). To ensure that the interviewees considered only the factors relevant to this study and to eliminate other influencing factors, the interviews began with the following statement: ‘Imagine you know nothing about a book, its author and its publisher, and then answer the following questions’.

Using Zoom software, the Arabic interviews were transcribed verbatim after each session, ensuring comparability, reliability, and consistency [18]. The duration of each interview ranged from fifteen to twenty-five minutes. The Arabic transcriptions were carefully translated into English and then imported into NVivo qualitative data analysis software for thematic analysis.

The qualitative interview questions were developed based on the quantitative data, and thematic analysis with its six-step process, as outlined by Braun and Clarke [19], was employed. The thematic analysis was conducted inductively and deductively. Each question represented a theme (i.e. OCR factor and related cultural characteristics) and was followed by

explanatory questions such as ‘Why?’ or statements such as ‘Please explain further’ to obtain deeper insights and uncover new factors related to participants’ OCR usage.

IV. RESULTS

This section presents the demographic statistics of the samples from the quantitative and qualitative phases of this study. Next, the steps of the statistical analysis to test the research hypotheses using PLS-SEM are discussed. Following that, the findings derived from the thematic analysis of the qualitative interviews are reported.

A. Sample profiles of quantitative and qualitative phases

During the quantitative stage of the research, 633 legitimate responses were gathered from Arab users of Goodreads; there were no invalid responses. Among them, 46% were male, while the majority (54%) were female. The vast majority of respondents (83%) were young adults 18–35 years old. Regarding their education level, two-thirds of them had a bachelor’s degree. Most of the respondents were from Saudi Arabia and Egypt.

The demographics of the sample align closely with the overall Goodreads audience statistics as analysed by Similarweb.com, a popular platform that provides insights and data about website traffic, audience demographics and other key metrics. Goodreads’ latest statistics reveal women constitute the majority (>60%) of the website audience, individuals younger than 35 years form the largest age group of visitors, and Egypt and Saudi Arabia rank among the top 30 countries for Goodreads.com traffic [20]. Thus, we can argue that the sample was representative of Arab users of Goodreads, and therefore the findings can be generalised across this population.

Regarding the participants’ experience with the OCRs, most of the respondents have long experience with using OCRs to aid their purchase decisions, with over two-thirds of them using OCRs for more than four years. Moreover, the participants reported checking OCRs very frequently before making purchases, with 30% always doing so and 40% often doing so. They typically spend a short amount of time reading reviews, with 54% spending less than 15 minutes. Furthermore, OCRs about both fiction and non-fiction books are important to the respondents, with almost two-thirds often reading reviews about both genres, whereas only 19% about fiction books and 17% about nonfiction books.

Regarding the qualitative research phase, fifteen online interviews with participants from the quantitative stage were conducted. Eight men and seven women comprised the sample, with the majority falling within the 26–35 age group. Different useful perspectives and explanations for the questionnaire data are presented in Section V.

B. Common method bias

When data are obtained from the same population at a single point in time, the validity of the study may be affected if the issue of common method bias exists [21]. Following the approach of most previous studies that addressed this issue

(e.g. [22], [23]), Harman's single-factor test was employed in SPSS. It involves loading all the items into a single factor. The result shows that the first factor explains only 25.5% of the data variation, well below the 50% level recommended by [21]. Therefore, the findings demonstrated that this study is not affected by the common method bias.

C. Measurement model

The initial stage of structural equation modelling (SEM) is examining the measurement models to verify the reliability and validity of the survey constructs. The internal consistency (reliability) is assessed by evaluating the value of Cronbach's alpha (α) and composite reliability (CR), while convergent validity is verified by examining the factor outer loadings and average variance extracted (AVE). The measurement model should meet all criteria as recommended by Hair et al [16] before the structural model can be assessed. First, indicator loadings are recommended to be above 0.7. Second, to establish construct reliability, α and CR values should be above 0.7, and for convergent validity, AVE value should be above 0.5. All results exceeded the thresholds except for loadings of two items (Img2 = 0.559 and Exp1 = 0.675), but loadings above 0.5 are acceptable if construct reliability and convergent validity are established [24], [25]. Thus, we retained them because convergent validity and internal reliability were achieved (see Table I).

Third, to establish discriminant validity, the square root of AVE for each construct should surpass the highest squared correlation with any other construct in the model (Fornell-Larcker criterion) [26]. Furthermore, Henseler et al. [27] developed a new and more accurate measurement for discriminant validity called the heterotrait-monotrait (HTMT) ratio of the correlations. The HTMT ratio should be below 0.85 [27]. As Table II illustrates, Fornell-Larcker and HTMT criteria were met, indicating that discriminant validity was achieved.

TABLE I.

CONSTRUCTS' RELIABILITY AND CONVERGENT VALIDITY.

Constructs	α	CR	AVE
Valence (Val)	0.793	0.866	0.617
Comprehensiveness (Com)	0.811	0.876	0.639
Readability (Read)	0.808	0.887	0.723
Image (Img)	0.831	0.873	0.638
Volume (Vol)	0.820	0.881	0.649
Experience (Exp)	0.811	0.876	0.641
Reputation (Rep)	0.873	0.913	0.725
Identity disclosure (ID)	0.901	0.931	0.772
Purchase intention (PI)	0.840	0.903	0.757

TABLE II.

DISCRIMINANT VALIDITY

	Val	Com	Read	Img	Vol	Exp	Rep	ID	PI
Val	0.786	0.463	0.421	0.051	0.494	0.462	0.396	0.403	0.490
Com	0.371	0.799	0.512	0.050	0.425	0.515	0.349	0.308	0.522

Read	0.338	0.418	0.850	0.073	0.445	0.562	0.417	0.353	0.461
Img	0.017	0.042	0.053	0.799	0.077	0.084	0.048	0.058	0.107
Vol	0.402	0.349	0.362	0.057	0.806	0.503	0.472	0.478	0.486
Exp	0.378	0.418	0.455	0.042	0.413	0.800	0.498	0.502	0.497
Rep	0.327	0.292	0.352	0.014	0.401	0.423	0.852	0.469	0.323
ID	0.338	0.264	0.301	0.002	0.413	0.432	0.424	0.878	0.301
PI	0.405	0.436	0.383	0.112	0.410	0.421	0.279	0.267	0.870

Underlined values = Fornell and Larcker criterion; italic values = heterotrait-monotrait (HTMT) criterion.

D. Multicollinearity

This is an important step for ensuring that the model constructs are not strongly correlated with each other. Values of the variance inflation factor (VIF) statistic above 5 indicate critical collinearity issues [16]. The results showed that the VIF values of the inner model range between 1.01 and 1.63, which is less than the conservative threshold of 3 [16]. Thus, multicollinearity is not a concern in this study.

E. Structural model

Having established an adequate measurement model, the structural model will be evaluated to test the research hypotheses. Using SmartPLS software, a bootstrapping approach with 5000 subsamples was used to address the significance (p-values) of the relationships between the research factors in the model [16]. As shown in Table III, review valence ($p < 0.01$), comprehensiveness ($p < 0.001$), readability ($p < 0.05$), image ($p < 0.05$), volume ($p < 0.001$) and reviewer experience ($p < 0.01$) positively and significantly impact the purchase intention. By contrast, reviewer identity disclosure and reputation were not statistically significant. Hence, H1, H2, H3, H4, H5 and H7 were accepted while H6 and H8 were rejected. The path coefficient, T-value and p-value for each research hypothesis are listed in Table III.

TABLE III.

HYPOTHESIS TESTING RESULTS

Hypothesis	β	T values	P values	Decision
H1: Val -> PI	0.169	3.344	0.001	Accepted
H2: Com -> PI	0.205	4.617	0.000	Accepted
H3: Read -> PI	0.109	2.250	0.025	Accepted
H4: Img -> PI	0.078	2.194	0.028	Accepted
H5: Vol -> PI	0.169	3.534	0.000	Accepted
H6: ID -> PI	-0.012	0.277	0.781	Rejected
H7: Exp -> PI	0.155	3.251	0.001	Accepted
H8: Rep -> PI	-0.004	0.072	0.943	Rejected

Next, the endogenous construct's coefficient of determination (R^2) value was calculated to assess the model's explanatory power. The model explains 33.8% of the variance in purchase intention, exceeding the benchmark often observed in customer decision-making research [28], [29]. In other words, review valence, comprehensiveness, readability, images, volume and reviewer experience can explain around one-third of

the change in the purchase intention. Fig 2 shows the revised research model with the significant paths.

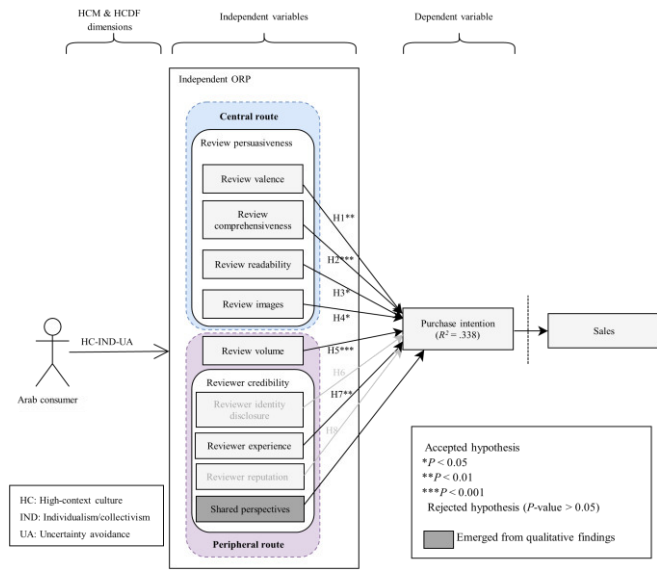


Fig 2. The revised AOCR-PI model

Finally, the value of cross-validated redundancy (Q2) determines the predictive relevance of the model, which estimates how well the model matches expectations. For the model to have predictive power, the Q2 value should be greater than zero [16]. The result shows that the Q2 value of the dependent construct (purchase intention) is 0.313. Hence, the model demonstrates predictive power.

F. Qualitative phase results

This section presents the thematic analysis findings from the qualitative interviews. Before delving into the main questions related to the OCR factors, the interviewees were asked about their usual usage of OCRs, specifically whether they prefer independent online review platforms (IORPs) such as Goodreads (external OCRs) or e-commerce websites such as Amazon (internal OCRs). As expected, most of them expressed a preference for IORPs. This question was followed by asking ‘Why?’ to gain insights into the reasons behind their preference for IORPs over e-commerce websites. The interview questions are available from the corresponding author upon request.

The following are the seven reasons reported by the interviewees for preferring external over internal OCRs. First, the most frequent reason revealed by the interviewees is trusting the reviews posted on IORPs, indicating that external OCRs are considered more reliable than internal OCRs. Second, the reviewers on IORPs are perceived as more credible by respondents. Third, the participants emphasised the widespread availability of OCRs on Goodreads compared to e-bookstores, which enables them to assess others’ opinions about books. Fourth, the reviewers on IORPs can freely provide their opinions without attempting to promote books. Fifth, IORPs offer better social features that allow users to follow each other and send private messages, indicating the

importance of social ties in Arab societies (HC culture). Sixth, Goodreads provides a feature called ‘reading challenge’ that allows users to set reading targets for the next twelve months, which was found to be influential according to the interviewed participants. Finally, some participants believe that external OCRs on IORPs provide a deeper evaluation of books compared to those found on e-commerce websites, which tend to be more superficial.

Following our deductive and inductive thematic analysis approach, we initially (deductively) coded interview responses to the pre-defined OCR factors from the research model (Fig 1). Inductive coding was also used to capture additional OCR-related factors not covered by the study framework. Information reported by at least two interviewees was classified under the ‘other factors’ theme. One code was found and added to the research model as represented in Fig 2 that is ‘shared perspectives’ between OCR users and reviewers. Fig 3 shows the thematic analysis findings, including themes related to OCR factors and associated child codes.

Moreover, while coding the transcripts, texts related to the cultural characteristics of Arabs (HC, IND and UA) were highlighted. Many words and phrases were found that align with HCM and HCDF. The next section discusses these findings in detail.

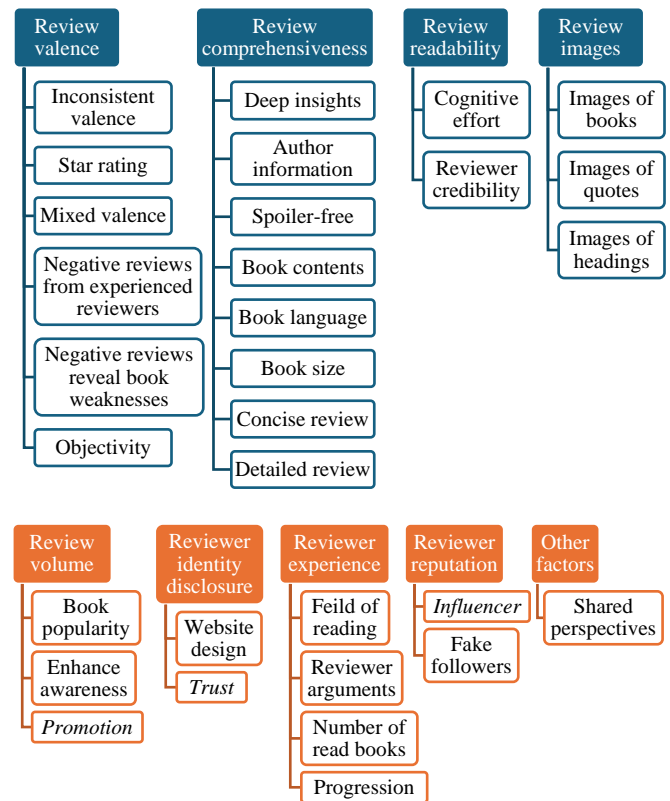


Fig 3. Thematic analysis results (text in italics indicates opposite viewpoints to the quantitative results)

V. DISCUSSION

According to prior research cultural backgrounds play a significant role in consumer decisions and Internet usage [30], [31], [32]. Drawing on ELM, HCM and HCDF, this study reports the effect of OCRs on the purchase intention of Arab book readers. Using PLS-SEM, eight hypotheses were tested. The study found that factors representing the content of the OCRs significantly influenced the study participants (H1, H2, H3 and H4). These factors represent the central cues of the review and require high cognitive effort from the reader. Review volume (H5) and reviewer experience (H7), which represent peripheral cues, also played significant roles in consumer purchase intention. Hence, the study asserts that to inform their buying decisions, Arab consumers primarily focus on the review content, particularly in the context of book reviews. This implies that ELM is an appropriate theoretical foundation for understanding how consumers use OCRs to aid their purchasing decisions. These findings are broadly supported by earlier research works (e.g. [31], [33], [34], [35]). A more detailed discussion of the findings is provided in the next sections.

A. Review-related Factors

The first central route factor investigated in this study is review valence. The results show that review valence was a significant construct that influenced Arabs' purchase intention of books. This finding aligns with previous studies (e.g. [36], [37], [38]) and can be attributed to the tendency of people from high UA cultures to be more emotional [39]. The valence of information is significantly influenced by emotional strength [40], meaning that positive and negative reviews often contain emotionally charged words. Because Arabs are strongly uncertainty-avoidant, their reviews may also be of high valence, consequently extreme OCRs could elicit emotional responses from readers and significantly affect the purchase intention of consumers. Therefore, consumers in this context value intense positive emotions as they seek reassurance and support before making product purchases.

The qualitative phase of this study revealed that both positive and negative reviews affect participants' purchase intentions. Interviewees indicated that they consider purchasing a book when there is a higher number of positive reviews compared to negative ones, which can be quickly determined from the overall rating. In addition, some interviewees revealed that the impact of review valence (H1) becomes more prominent when the reviewer has more experience (H7) with the book topic. This is supported by Casaló et al. [41], who found that reviewer expertise increases the value of positive OCRs. Another point mentioned by some interviewees was that reviewers should include details and reasons (H2) behind the positive and negative reviews, rather than providing very short reviews such as 'the book is well-translated' or 'bad book'. Therefore, consumers' likelihood of making a purchase is significantly influenced if positive/negative OCRs are informative (comprehensive) and posted by trustworthy (credible) users [42]. Some interviewees use star ratings as an easily

interpreted sign of valence, and this finding is consistent with previous studies which used them to represent valence [43]. Thus, this study confirms that star ratings can be used as an indicator of review valence.

The results reveal the crucial role of comprehensiveness and volume of reviews in shaping the purchase intention of OCR receivers. Their importance is likely due to the fact that detailed reviews provide much information, enabling receivers to evaluate products precisely. Hence, a comprehensive review is perceived as valuable and persuasive. This aligns with Bosman et al. [44] who found that users consider the length of book reviews a good measure of review credibility, thus enhancing the effectiveness of their purchasing decisions. This impact is particularly significant for Arab consumers since they are more inclined to avoid uncertainty [45] and therefore seek to reduce ambiguity by obtaining additional information. Furthermore, because books are experience products, readers require considerable information to evaluate their quality. Additionally, as the qualitative data showed, consumers interpreted a substantial quantity of reviews as a sign of product popularity, consequently enhancing their purchase intentions. This finding aligns with Zhao et al. [46] who state that "an increase in the number of a product's online reviews will result in an increase in customers' intention to initiate the purchase of that product". Their study was conducted with Chinese consumers, who share characteristics with Arabs, such as being HC, collectivistic and uncertainty-avoidant. Therefore, the comprehensiveness and quantity of reviews significantly influence the purchasing intention of Arab customers by providing valuable information, reducing ambiguity and indicating product popularity.

Another finding is the persuasive impact of images on review readers. A similar result was reported regarding cultures similar to the Arab culture. For instance, Cheng and Ho [47] found that photos of dining (an experience good) are very important to Chinese consumers during the decision-making process because images provide additional information (i.e. improve argument quality). In the qualitative phase of this study, most interviewees stated that images such as book covers and contents affect their purchase intention. Therefore, we argue that when both textual and visual information is provided, OCR receivers can have a more complete knowledge of a product, which reduces their uncertainty and aids them in their purchasing decisions. This result adds to previous studies that confirmed the impact of images posted by customers on individual intentions to visit destinations [48], [49]. Moreover, it confirms that Arabs, an HC culture, prefer visual information, which significantly influences their purchase intentions.

Moreover, the results show that Arab consumers prefer easy-to-read and understandable OCRs. The qualitative findings provide several explanations for this result. One interviewee stated that a readable review written in formal language "indicates the maturity of the reader [reviewer] and their deeper understanding" of the topic of the reviewed book.

This explanation indicates that OCR receivers can identify experienced reviewers by the way they write their reviews. Similarly, another interviewee believes that a review that is error-free ‘gives the impression that the reviewer is someone whose words can be trusted’. In addition, the readability of the text has an impact on how reviews are perceived as valuable [50] and consequently, the purchase intention of consumers. With respect to the ELM, simple text requires less cognitive effort and is therefore more useful than complicated text [51]. Hence, consumers can evaluate the product quality and make purchase decisions more easily and quickly from a readable OCR.

B. Reviewer-related Factors

The reviewer credibility factors are represented by the reviewer identity disclosure, experience and reputation. These act as peripheral cues that define the believability of the information. The results reveal that Arab users of book ORPs infrequently check reviewers’ profile information and pattern factors to rationalise their buying decisions. Experience was found to be the most important reviewer characteristic affecting consumers’ intention to purchase books. In other words, the more reviews provided by a reviewer, the higher their credibility [52] and the more effective the message delivered to the receiver [53], which, in turn, affected the consumer purchase intention. This finding is supported by Baek et al. [54] who indicated that reviews originating from an experienced source are more meaningful in the minds of readers and have a more significant influence compared with reviews from sources without expertise.

Interestingly, this study found that reviewer identity disclosure and reputation, as defined by the number of followers and friends, have no impact on Arabs reading OCRs to aid their purchasing intentions. This finding is consistent with the research conducted by Baek et al. [54], which demonstrated that the reviews provided by top-ranked reviewers exert a substantial impact on the helpfulness of the review, but disclosing the reviewer’s real name does not. However, most current research illustrates that reviewer identity and reputation have a significant impact on consumer behaviour (e.g. [55], [56]). These inconsistent results may be explained by the fact that most previous studies were conducted on e-commerce websites, whereas this research focused on Goodreads, a third-party ORP. Furthermore, to the best of our knowledge, none of them studied the impact of these reviewer-related factors on Arab OCR users.

The qualitative findings offer insight into the unexpected quantitative results regarding the impact of the reviewer identity disclosure and reputation on purchase intention. For example, one interviewee mentioned not being concerned about the number of followers, as reviewers can purchase fake followers to be in the ‘most followed’ list. Another interviewee argued that some reviewers have many followers not due to the quality of their reviews but because they are celebrities or well-known book authors. The non-significant impact of the personal reviewer information on purchase intention can be

explained by the concept of egocentric anchoring, which suggests that when consumers have insufficient information about the reviewer, they tend to mentally fill in the gaps in the reviewers’ identities with their own assumptions [57]. As a result, consumers are more likely to be influenced by reviews written by anonymous reviewers compared with reviews from dissimilar reviewers [57]. This explanation is further supported by the emergent qualitative theme of similarity between the interests of the reviewer and the OCR receiver. Although the reviewer’s identity and reputation are not significant, over half of those interviewed emphasised the importance of checking the reviewers’ profiles. Therefore, this research asserts that book reviewers should make their profiles public, allowing ORP users to access their previous reviews and book genres.

The qualitative study encouraged participants’ freedom of expression, revealing an important factor: similarity in perspectives between reviewers and OCR recipients. When reviewers share similar interests and viewpoints with the OCR recipients, their credibility increases, and their reviews have a more pronounced impact on purchase intention compared to those expressing contrasting opinions. This concept is known as ‘homophily’ in previous research, defined as “the degree to which two or more individuals who interact are similar in certain attributes” [58]. Homophily typically relates to similarity in demographic characteristics or perceived attributes (preferences and values) [53]. However, in this study, we introduce the term ‘shared perspectives’ to emphasise the homophily in perceived attributes between the OCR sender and receiver, given that the demographic information was found to be non-significant, as discussed earlier. Therefore, this factor was added later to the AOOCR-PI model, as shown in Fig 2.

This finding aligns with HCM and HCDF theories and previous studies. For instance, since HC cultures (e.g. Arabs) have a high tendency to use implicit cues and shared common knowledge [9], the similarity of perspectives between the reviewer and the OCR receiver leverages this cultural preference for shared understanding. Furthermore, as a collectivist culture, Arabs emphasise cohesive in-groups and the opinions of others [39]. This finding is in agreement with Cheong and Mohammed-Baksh’s [59] research, which demonstrated the significant influence of reference group recommendations on the purchase decisions of South Korean consumers. They further suggested that their findings are potentially applicable to other countries with cultural similarities to South Korea, particularly those that exhibit HC cultures, such as Arabic culture. Therefore, this finding highlights the importance of considering this when designing online review systems.

C. Theoretical Contribution

This research provides a new understanding of OCR usage to aid the purchase intention of Arab book readers. The findings from this study make several theoretical contributions to the current literature. First, this work contributes to the ELM by providing empirical evidence that supports the validity of the model in the context of the OCR research field. The ELM

suggests that individuals need high cognitive effort when processing messages (i.e. use the central route), whereas when they lack interest or motivation towards the message or encounter challenges in comprehending arguments pertaining to the relevant issue, they use the peripheral route [60]. This study found that Arab users of OCRs start by seeking the needed information from the textual reviews while the quantity of reviews and reviewer information act as peripheral cues that aid in their purchasing decision-making. Moreover, the qualitative results illustrate the interaction effect between central and peripheral routes. For instance, OCR receivers use review readability and comprehensibility as indicators of a reviewer's experience. This provides new evidence that the central and peripheral routes complement each other and work simultaneously. Therefore, this study confirms that the ELM is a suitable theoretical foundation for exploring the impact of OCRs on purchase intention.

Second, this research applies HCM and HCDM theories to explain Arab consumer behaviour in the online environment. Although these theories were initially developed to describe face-to-face communication, this study confirms their relevance in explaining the impact of OCRs on consumer purchase intention. The main findings of this paper support Hofstede's [10] and some aspects of Hall's [9] theories regarding OCR usage by Arabs. Specifically, the results affirm the collectivism and uncertainty-avoidant nature of Arab consumers, indicating that they are influenced by the collective consensus when forming purchasing decisions. In addition, Arab book consumers exhibit a marked preference for extensive information gathering, through OCRs, to mitigate perceived risk associated with future purchases. Furthermore, this research reveals that the previous activity of a reviewer (experience) significantly influences Arab OCR users and that is in line with HCM, which states that Arab individuals may take a longer time to build relationships with others [61]. However, direct and explicit language is preferred by Arab users of OCRs, at least in the context of book readers, although Arab societies tend to use indirect and implicit messaging [9]. These results may also apply to other countries with HC cultures, such as Asian, Eastern European and Latin countries.

Third, this study reveals an interesting finding regarding the impact of reviewer identity and reputation on purchase intention. In contrast to previous studies which emphasised the significance of a reviewer's personality and popularity on consumer decisions (e.g. [56], [62]), the findings of this research indicate that Arab consumers prioritise the experience of the reviewer over their identity and number of followers. However, it is important to note that this outcome may not be applicable to other cultures and products since, to the best of our knowledge, this study is the first of its type conducted on book reviews in the Arabic context. To confirm or contradict these findings, further work is needed to address the impact of reviewer-related attributes on purchase intention, accounting for varying types of products and ORPs.

D. Practical Implications

This study provides valuable insights for ORP developers, book authors and publishers seeking to understand the factors influencing Arab consumers when purchasing books. The outcomes of this study could be employed to enhance online review systems that align with Arab preferences, ultimately leading to increased book sales.

Considering that current online review systems are predominantly designed in the West, the findings can guide web developers of Arabic ORPs in building or customising current systems by prioritising the reviews based on the most important factors for Arab users. For example, rather than simply sorting reviews by posting date, developers may employ natural language processing methods, such as sentiment analysis, and develop new machine learning algorithms that consider review readability and comprehensiveness. Therefore, the outcomes of this study can assist researchers and professionals in minimising cultural bias when creating, evaluating and validating algorithms that rely on OCRs.

Moreover, the outcomes of this study have implications for managing OCRs and reducing consumer uncertainty. Administrators of e-commerce and publisher websites should refrain from editing or removing negative reviews but rather encourage reviewers to share more information freely, which enables authors and publishers to identify weaknesses and work towards satisfying readers in the next edition. This is because a large number of reviews can enhance consumer confidence, which, in turn, can reduce consumer uncertainty and increase purchase intentions [35]. However, it is important to note that the abundance of OCRs available to consumers may cause information overload [50], [63]. Therefore, these results need to be interpreted with caution, and further research is required to determine the optimal number of reviews required for Arab consumers to make informed purchase decisions.

VI. CONCLUSION

The present study was designed to determine the effect of OCRs on the purchase intention of Arab consumers. To achieve this objective, a mixed-methods research design was adopted, consisting of an exploratory quantitative study followed by an explanatory qualitative study. The data were collected from Arab users of Goodreads. The study revealed that customers place great importance on the central cues related to the review content, namely, the review valence, comprehensiveness, images and readability. In addition, the empirical data analysis indicated that the review volume (number of reviews) and reviewer experience (number of reviews provided by a reviewer) work as peripheral cues affecting the purchase intention of Arab book readers. Furthermore, the qualitative results showed that shared perspectives between reviewers and OCR receivers significantly impact the purchase intention of Arab consumers. However, the reviewer's identity and number of friends and followers (reputation) were found to be non-significant factors for Arabs when using OCRs for book purchases. The findings that emerged from

this research offer useful insights for ORPs and recommendation system developers as well as book authors and publishers on how to enhance the experience of OCR users, encourage book purchases and consequently increase sales.

Several important limitations need to be considered, and the results must be interpreted cautiously. Nevertheless, these limitations provide insights into new opportunities for further research. First, this research was conducted on Arabs from various countries, with data collected exclusively from Goodreads users. The findings of the study might have varied if the participants' nationalities were considered, as there may be cultural differences between Arab countries, as indicated by Hofstede's dimensions. Additionally, to test the model with different audiences, it would be interesting to collect and compare data from different ORPs.

Second, this study focused solely on OCRs of books, which are considered experience products. Therefore, the results may not be directly applicable to search products. To broaden the scope, it is recommended that further research be undertaken on other types of goods.

Third, the dependent variable, purchase intention, is a weaker predictor of book sales compared with actual purchases. However, because sales data are difficult to access and Goodreads is the largest and most popular online platform for book reviews, with over 120 million users [64], we argue that purchase intention is an appropriate, available and easily measured predictor. Nevertheless, it is recommended that future research use actual sales data and compare the results with the findings of this study.

Finally, due to time constraints and to limit the survey completion time, this study investigated the most important OCR factors reported in the literature. However, further research could explore other review and reviewer factors, such as the similarity in interests and perspectives between the reviewer and the OCR user. These two factors represent the peripheral route in the ELM as they are not directly related to the content of the review. Although these two factors have been studied extensively, insufficient attention has been paid to their influence in the Arabic context.

REFERENCES

- [1] "Survey: The Ever-Growing Power of Reviews (2023 Edition)." Accessed: May 06, 2024. [Online]. Available: <https://www.powerreviews.com/power-of-reviews-2023/>
- [2] X. Li, C. Wu, and F. Mai, "The effect of online reviews on product sales: A joint sentiment-topic analysis," *Information and Management*, vol. 56, no. 2, pp. 172–184, 2019, doi: 10.1016/j.im.2018.04.007.
- [3] S. M. Mudambi and D. Schuff, "What makes a helpful review? A study of customer reviews on amazon.com," *MIS Quarterly*, vol. 34, no. 1, pp. 185–200, 2010, doi: 10.1016/j.ica.2011.08.067.
- [4] H. Hong, D. Xu, G. A. Wang, and W. Fan, "Understanding the determinants of online review helpfulness: A meta-analytic investigation," *Decision Support Systems*, vol. 102, pp. 1–11, 2017, doi: 10.1016/j.dss.2017.06.007.
- [5] B. Gu, J. Park, and P. Konana, "The impact of external word-of-mouth sources on retailer sales of high-involvement products," *Information Systems Research*, vol. 23, no. 1, pp. 182–196, 2012, doi: 10.1287/isre.1100.0343.
- [6] C. Luo, J. Wu, Y. Shi, and Y. Xu, "The effects of individualism–collectivism cultural orientation on eWOM information," *International Journal of Information Management*, vol. 34, pp. 446–456, 2014, doi: 10.1016/j.ijinfomgt.2014.04.001.
- [7] J. E. M. Steenkamp and I. Geyskens, "How Country Characteristics Affect the Perceived Value of Web Sites," *Journal of Marketing*, vol. 70, no. 3, pp. 136–150, 2006, doi: 10.1509/jmkg.70.3.136.
- [8] R. E. Petty and J. T. Cacioppo, "The elaboration likelihood model of persuasion," in *Advances in Experimental Social Psychology*, vol. 19, no. C, 1986, pp. 123–205. doi: 10.1016/S0065-2601(08)60214-2.
- [9] E. T. Hall, *Beyond culture*. Garden City, N.Y.: Anchor Press, 1976. doi: 10.4324/9780203894880.ch3.
- [10] G. Hofstede, "Culture and Organizations," *International Studies of Management & Organization*, vol. 10, no. 4, pp. 15–41, Dec. 1980, doi: 10.1080/00208825.1980.11656300.
- [11] A. Alghamdi, N. Beloff, and M. White, "A New Arabic Online Consumer Reviews Model to Aid Purchasing Intention (AOCR-PI)," in *Lecture Notes in Networks and Systems*, K. (eds) I. S. and Applications. I. 2022 Arai, Ed., Cham: Springer, 2023, pp. 475–492. doi: 10.1007/978-3-031-16072-1_35.
- [12] W. Kai, L. Xiaojuan, and H. Yutong, "Exploring Goodreads reviews for book impact assessment," *Journal of Informetrics*, vol. 13, pp. 874–886, 2019, doi: 10.1016/j.joi.2019.07.003.
- [13] R. E. Petty, J. T. Cacioppo, A. J. Strathman, and J. R. Priester, "To think or not to think. Exploring two routes to persuasion," in *Persuasion: Psychological insights and perspectives*, 2005, pp. 81–116.
- [14] J. W. Creswell and V. L. P. Clark, *Designing and Conducting Mixed Methods Research*, 3rd ed., no. 1. Thousand Oaks, California: SAGE Publications, Inc, 2018. doi: 10.1177/1937586719832223.
- [15] J. Berger, "Word of mouth and interpersonal communication: A review and directions for future research," *Journal of Consumer Psychology*, vol. 24, no. 4, pp. 586–607, 2014, doi: 10.1016/j.jcps.2014.05.002.
- [16] J. F. Hair, J. J. Risher, and C. M. Ringle, "When to use and how to report the results of PLS-SEM," vol. 31, no. 1, pp. 2–24, 2019, doi: 10.1108/EBR-11-2018-0203.
- [17] M. N. K. Saunders, P. Lewis, and A. Thornhill, *Research methods for business students*, 6th ed. Harlow, England ; New York: Pearson, 2012.
- [18] D. Silverman, *Doing Qualitative Research*, 5th ed. London: SAGE Publications Ltd, 2017.

- [19] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006, doi: 10.1191/1478088706qp0630a.
- [20] "goodreads.com Traffic Analytics, Ranking & Audience [March 2024]," Similarweb. Accessed: May 06, 2024. [Online]. Available: <https://www.similarweb.com/website/goodreads.com/>
- [21] P. M. Podsakoff, S. B. MacKenzie, J. Y. Lee, and N. P. Podsakoff, "Common Method Biases in Behavioral Research: A Critical Review of the Literature and Recommended Remedies," *Journal of Applied Psychology*, vol. 88, no. 5, pp. 879–903, 2003, doi: 10.1037/0021-9010.88.5.879.
- [22] A. Sharma, Y. K. Dwivedi, V. Arya, and M. Q. Siddiqui, "Does SMS advertising still have relevance to increase consumer purchase intention? A hybrid PLS-SEM-neural network modelling approach," *Computers in Human Behavior*, vol. 124, no. January, p. 106919, 2021, doi: 10.1016/j.chb.2021.106919.
- [23] Z. Sheikh, T. Islam, S. Rana, Z. Hameed, and U. Saeed, "Acceptance of social commerce framework in Saudi Arabia," *Telematics and Informatics*, vol. 34, no. 8, pp. 1693–1708, 2017, doi: 10.1016/j.tele.2017.08.003.
- [24] J. Hair Jr, G. Hult, C. Ringle, and M. Sarstedt, *A Primer on Partial Least Squares Structural Equation Modeling (PLS-SEM)*. 2017.
- [25] J. Henseler, C. M. Ringle, and R. R. Sinkovics, "The use of partial least squares path modeling in international marketing," *Advances in International Marketing*, vol. 20, pp. 277–319, 2009, doi: 10.1108/S1474-7979(2009)0000020014.
- [26] C. Fornell and D. F. Larcker, "Evaluating Structural Equation Models with Unobservable Variables and Measurement Error," *Journal of Marketing Research*, vol. 18, no. 1, p. 39, Feb. 1981, doi: 10.2307/3151312.
- [27] J. Henseler, C. M. Ringle, and M. Sarstedt, "A new criterion for assessing discriminant validity in variance-based structural equation modeling," *Journal of the Academy of Marketing Science*, vol. 43, no. 1, pp. 115–135, 2015, doi: 10.1007/s11747-014-0403-8.
- [28] A. Shukla and A. Mishra, "Role of Review Length, Review Valence and Review Credibility on Consumer's Online Hotel Booking Intention," *FIIB Business Review*, vol. 12, no. 4, pp. 403–414, 2022, doi: 10.1177/23197145221099683.
- [29] A. Shukla and A. Mishra, "Effects of Visual Information and Argument Concreteness on Purchase Intention of Consumers Towards Online Hotel Booking," *Vision*, vol. 27, no. 5, pp. 639–649, 2021, doi: 10.1177/097226292111038069.
- [30] R. A. L. Fischer, R. Walczuch, and E. Guzman, "Does Culture Matter? Impact of Individualism and Uncertainty Avoidance on App Reviews," *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*, pp. 67–76, 2021, doi: 10.1109/ICSE-SEIS52602.2021.00016.
- [31] M. R. González-Rodríguez, M. C. Díaz-Fernández, A. Bilgihan, F. Okumus, and F. Shi, "The impact of eWOM source credibility on destination visit intention and online involvement: a case of Chinese tourists," *Journal of Hospitality and Tourism Technology*, vol. 13, no. 5, pp. 855–874, 2022, doi: 10.1108/JHTT-11-2021-0321.
- [32] A. ur Rehman, "Consumers' perceived value of luxury goods through the lens of Hofstede cultural dimensions: A cross-cultural study," *Journal of Public Affairs*, vol. 22, no. 4, 2021, doi: 10.1002/pa.2660.
- [33] Y. Li, X. Wang, and C. Van Slyke, "Determinants of online professor reviews: an elaboration likelihood model perspective," *Internet Research*, 2022, doi: 10.1108/INTR-11-2020-0627.
- [34] S. G. Moore and K. C. Laffreniere, "How online word-of-mouth impacts receivers," *Consumer Psychology Review*, vol. 3, no. 1, pp. 34–59, 2020, doi: 10.1002/arcp.1055.
- [35] S. Teng, K. W. Khong, A. Y. L. Chong, and B. Lin, "Examining the impacts of electronic word-of-mouth message on consumers' attitude," *Journal of Computer Information Systems*, vol. 57, no. 3, pp. 238–251, 2017, doi: 10.1080/08874417.2016.1184012.
- [36] J. W. Cheong, S. Muthaly, M. Kuppasamy, and C. Han, "The study of online reviews and its relationship to online purchase intention for electronic products among the millennials in Malaysia," *Asia Pacific Journal of Marketing and Logistics*, vol. 32, no. 7, pp. 1519–1538, 2020, doi: 10.1108/APJML-03-2019-0192.
- [37] E. Ismagilova, E. L. Slade, N. P. Rana, and Y. K. Dwivedi, "The Effect of Electronic Word of Mouth Communications on Intention to Buy: A Meta-Analysis," *Information Systems Frontiers*, vol. 22, no. 5, pp. 1203–1226, 2019, doi: 10.1007/s10796-019-09924-y.
- [38] M. A. Sutanto and A. Aprianingsih, "the Effect of Online Consumer Review Toward Purchase Intention: a Study in Premium Cosmetic in Indonesia," *International Conference on Ethics of Business, Economics, and Social Science*, pp. 218–230, 2016.
- [39] G. Hofstede, "Dimensionalizing Cultures: The Hofstede Model in Context," *Online Readings in Psychology and Culture*, vol. 2, no. 1, pp. 2307–0919, 2011, doi: 10.9707/2307-0919.1014.
- [40] J. Li and L. Zhan, "How the written word drives WOM: Evidence from consumer-generated product reviews," *Journal of Advertising Research*, vol. 51, no. 1, pp. 239–257, 2011, doi: 10.2501/JAR-51-1-239-257.
- [41] L. V. Casaló, C. Flavián, M. Guinalfú, and Y. Ekinci, "Avoiding the dark side of positive online consumer reviews: Enhancing reviews' usefulness for high risk-averse travelers," *Journal of Business Research*, vol. 68, no. 9, pp. 1829–1835, 2015, doi: 10.1016/j.jbusres.2015.01.010.

- [42] K. Z. K. Zhang, S. J. Zhao, C. M. K. Cheung, and M. K. O. Lee, "Examining the influence of online reviews on consumers' decision-making: A heuristic-systematic model," *Decision Support Systems*, vol. 67, pp. 78–89, 2014, doi: 10.1016/j.dss.2014.08.005.
- [43] K. L. Xie, Z. Zhang, and Z. Zhang, "The business value of online consumer reviews and management response to hotel performance," *International Journal of Hospitality Management*, vol. 43, pp. 1–12, 2014, doi: 10.1016/j.ijhm.2014.07.007.
- [44] D. J. Bosman, C. Boshoff, and G.-J. Van Rooyen, "The review credibility of electronic word-of-mouth communication on e-commerce platforms," *Management Dynamics*, vol. 22, no. 3, pp. 29–44, 2013.
- [45] G. Hofstede, G. J. Hofstede, and M. Minkov, *Cultures and organizations: software of the mind: intercultural cooperation and its importance for survival*, vol. 17, no. 4. New York: McGraw-Hill, 2010. doi: 10.1177/030630709201700409.
- [46] Z. Zhao, J. Wang, H. Sun, Y. Liu, Z. Fan, and F. Xuan, "What Factors Influence Online Product Sales? Online Reviews, Review System Curation, Online Promotional Marketing and Seller Guarantees Analysis," *IEEE Access*, vol. 8, pp. 3920–3931, 2020, doi: 10.1109/ACCESS.2019.2963047.
- [47] Y. H. Cheng and H. Y. Ho, "Social influence's impact on reader perceptions of online reviews," *Journal of Business Research*, vol. 68, no. 4, pp. 883–887, 2015, doi: 10.1016/j.jbusres.2014.11.046.
- [48] R. Filieri, Z. Lin, G. Pino, S. Alguezaui, and A. Inversini, "The role of visual cues in eWOM on consumers' behavioral intention and decisions," *Journal of Business Research*, vol. 135, no. June, pp. 663–675, 2021, doi: 10.1016/j.jbusres.2021.06.055.
- [49] T. Y. Wang and J. Park, "Destination Information Search in Social Media and Travel Intention of Generation Z University Students," *Journal of China Tourism Research*, vol. 19, no. 3, pp. 570–588, 2023, doi: 10.1080/19388160.2022.2101574.
- [50] B. Fang, Q. Ye, D. Kucukusta, and R. Law, "Analysis of the perceived value of online tourism reviews: Influence of readability and reviewer characteristics," *Tourism Management*, vol. 52, pp. 498–506, 2016, doi: 10.1016/j.tourman.2015.07.018.
- [51] A. Agnihotri and S. Bhattacharya, "Online Review Helpfulness: Role of Qualitative Factors," *Psychology & Marketing*, vol. 33, no. 11, pp. 1006–1017, 2016, doi: 10.1002/mar.
- [52] I. Syafganti and M. Walrave, "Assessing the Effects of Valence and Reviewers' Expertise on Consumers' Intention to Book and Recommend a Hotel," *International Journal of Hospitality and Tourism Administration*, vol. 23, no. 5, pp. 904–923, 2022, doi: 10.1080/15256480.2021.1881939.
- [53] B. Lis, "In eWOM We Trust: A Framework of Factors that Determine the eWOM Credibility," *Bus Inf Syst Eng*, vol. 5, no. 3, pp. 129–140, Jun. 2013, doi: 10.1007/s12599-013-0261-9.
- [54] H. Baek, J. Ahn, and Y. Choi, "Helpfulness of online consumer reviews: Readers' objectives and review cues," *International Journal of Electronic Commerce*, vol. 17, no. 2, pp. 99–126, 2012, doi: 10.2753/JEC1086-4415170204.
- [55] S. R. Hill, I. Troshani, and D. Chandrasekar, "Signalling Effects of Vlogger Popularity on Online Consumers," *Journal of Computer Information Systems*, vol. 60, no. 1, pp. 76–84, Jan. 2020, doi: 10.1080/08874417.2017.1400929.
- [56] J. Li and X. Liang, "Reviewers' Identity Cues in Online Product Reviews and Consumers' Purchase Intention," *Frontiers in Psychology*, vol. 12, Jan. 2022, doi: 10.3389/fpsyg.2021.784173.
- [57] R. W. Naylor, C. P. Lamberton, and D. A. Norton, "Seeing Ourselves in Others: Reviewer Ambiguity, Egocentric Anchoring, and Persuasion," *Journal of Marketing Research*, vol. 48, no. 3, pp. 617–631, 2011, doi: 10.1509/jmkr.48.3.617.
- [58] E. Ismagilova, E. Slade, N. P. Rana, and Y. K. Dwivedi, "The effect of characteristics of source credibility on consumer behaviour: A meta-analysis," *Journal of Retailing and Consumer Services*, vol. 53, p. 101736, Mar. 2020, doi: 10.1016/j.jretconser.2019.01.005.
- [59] H. J. Cheong and S. Mohammed-Baksh, "U.S. and Korean Consumers: A Cross-Cultural Examination of Product Information-Seeking and -Giving," *Journal of Promotion Management*, vol. 26, no. 6, pp. 893–910, 2020, doi: 10.1080/10496491.2020.1745985.
- [60] G. Roy, B. Datta, and R. Basu, "Effect of eWOM Valence on Online Retail Sales," *Global Business Review*, vol. 18, no. 1, pp. 198–209, 2017, doi: 10.1177/0972150916666966.
- [61] E. T. Hall and M. R. Hall, *Understanding cultural differences*. Yarmouth, ME: Intercultural Press, 1990. doi: 10.4324/9781315277349-2.
- [62] Z. Zhu, J. Liu, and W. Dong, "Factors correlated with the perceived usefulness of online reviews for consumers: a meta-analysis of the moderating effects of product type," *AJIM*, vol. 74, no. 2, pp. 265–288, Feb. 2022, doi: 10.1108/AJIM-02-2021-0054.
- [63] R. Zinko, P. Stolk, Z. Furner, and B. Almond, "A picture is worth a thousand words: how images influence information quality and information load in online reviews," *Electronic Markets*, vol. 30, no. 4, pp. 775–789, 2020, doi: 10.1007/s12525-019-00345-y.
- [64] M. Walsh and M. Antoniak, "The goodreads 'classics': A computational study of readers, amazon, and crowdsourced amateur criticism," *Journal of Cultural Analytics*, vol. 4, no. 1, pp. 243–287, 2021, doi: 10.22148/001C.22221.

Empirical Insights into Cloud Adoption: A new Model Exploring Influencing Factors for Saudi Arabian Small and Medium Enterprises

Mohammed Alqahtani
0009-0001-4915-4508
College of Computers and
Information Technology,
University of Bisha,
Saudi Arabia.
Department of Informatics,
University of Sussex, Brighton,
United Kingdom
ma2251@sussex.ac.uk

Natalia Beloff
0000-0002-8872-7786
Department of Informatics,
University of Sussex, Brighton,
United Kingdom
N.Beloff@sussex.ac.uk

Martin White
0000-0001-8686-2274
Department of Informatics,
University of Sussex, Brighton,
United Kingdom
m.white@sussex.ac.uk

Abstract—Cloud computing technology has emerged as a crucial driver of success for Small and Medium Enterprises (SMEs) globally, accelerating work processes and optimizing operations. Notably, SMEs in developed nations, including the United States and the United Kingdom, have proactively harnessed Cloud computing services, reaping substantial benefits in operational efficiency and time utilization. However, in many developing countries, including Saudi Arabia, most SMEs continue to rely on traditional technology, such as On-Premises Servers, instead of Cloud computing services. To investigate the factors influencing Cloud adoption, a new empirical model, the Adoption of Cloud Computing Model for Saudi Arabian SMEs (ACCM-SME), was developed. This study collected quantitative data from 412 participants representing Saudi SMEs in Riyadh city. The empirical data analysis revealed that 12 out of the 17 tested hypotheses exhibited significant positive influence, while five hypotheses failed to meet the specified research criteria and were consequently rejected. This research underscores the critical need to accelerate Cloud technology adoption among SMEs in developing countries, particularly Saudi Arabia. Bridging this technology gap has the potential to significantly enhance SMEs' competitiveness and operational efficiency, contributing to overall economic development. The ACCM-SME model provides nuanced insights into the factors influencing Cloud adoption, guiding further research. The study's rejected hypotheses highlight areas requiring attention for successful adoption. Policymakers and business leaders can leverage these findings to formulate strategies that facilitate Cloud adoption among SMEs.

Index Terms—Cloud computing adoption, Saudi SMEs, Cloud services.

I. INTRODUCTION

CLOUD computing has the potential significantly enhance small to medium enterprises' (SME) competitiveness and operational efficiency, contributing to overall economic development. Critical to adoption of Cloud technology by SMEs is better understanding of the influencing factors leading to SME adoption.

Cloud computing, in recent years, has been an effective technology that accelerates business and contributes to its success after the technological exploitation made in Internet services. Cloud Computing has been found to play a significant role in the success of Small and Medium Enterprises (SMEs) from different perspective. Growth in the

use of Cloud computing services among SMEs seeking new revenue opportunities has been fast over the previous five years, with a compound annual growth rate (CAGR) of 18% [1]. In addition, Cloud Computing Technology (CCT) has become a game-changer for improving inter-organisational collaboration [2]. What makes the Cloud these days so popular and targeted by many organisations worldwide is that users can clearly use Cloud computing services on demand immediately and at affordable prices without making any further investments in hardware or software updates [3]. In the last ten years, in many developed countries, including the United States and the United Kingdom, SMEs have Utilized and adopted Cloud computing technology services effectively and widely, which contributes to reducing their cost and processing their tasks continuously without any disruption [4], [5].

In contrast, in many developing countries, SMEs still face challenges and obstacles that slow their conversion to Cloud services due to some hindering factors. For instance, in Saudi Arabia, where this study was conducted, despite significant growth in internet services, mainly Cloud computing services, in recent years, dealing with sophisticated security challenges still needs more research and knowledge [6], [7]. According to a study released by the International Trade Administration, the Kingdom of Saudi Arabia is among the top nations in the Middle East regarding investing in and promoting information and communications technology (ICT)[8]. However, the adoption of Cloud computing services seems to have not received abundant attention until 2020, when the Ministry of Communications and Information Technology (MCIT) officially announced the strategy of Saudi Arabia to adopt Cloud computing.

This study aims to comprehensively examine the adoption of Cloud computing services by Saudi SMEs. The objectives encompass a thorough review of existing literature [9], to unveil the barriers and challenges influencing adoption, identification of the most impactful factors from the perspectives of top managers, IT managers, and employees, investigation into the effects of Cloud computing adoption on the mentioned stakeholders, and validation of a suitable

conceptual framework for a nuanced understanding of this adoption within the context of Saudi SMEs.

The study results show the benefits of shifting to Cloud services for SMEs. In addition, two limitations have been recognised and should be investigated further: potential job opportunities decreasing after adopting the Cloud and the perspective of Cloud providers. These limitations offer an opportunity for future research to study them and explore their effect on Cloud adoption.

In this study, the factors that hinder the adoption of Cloud computing services by Saudi SMEs have been investigated in the four primary contexts: Technological, Organisational, Environmental, and Social. This study aims to determine what factors affect Cloud adoption and which do not. Thus, a novel Framework was designed for this study to identify the factors impacting Saudi SMEs' adoption of Cloud computing services based on the TOE Model blended with DOI theory. This study subsequently evaluates the real effects of factors by analysing the collected data. Then, it will provide a valuable contribution towards enhancing the perception of Cloud computing, hence increasing the intention of SMEs in Saudi Arabia to adopt Cloud computing services.

This study aims to enhance the intention of Saudi SMEs to adopt Cloud computing services. Through an examination of Cloud adoption across four primary contexts, this research develops 17 hypotheses aimed at understanding and increasing Cloud adoption among Saudi SMEs. Thus, any accepting or rejection hypothesis will contribute to clearly showing the factors that will help Saudi SMEs increase their intention and direction toward adopting Cloud computing services.

The subsequent sections of this paper will sequentially unfold the hypotheses of the study, introduce the framework, and elucidate the methodology employed. The discussion on methodology will be complemented by an exploration of measurements and assessment tools, providing a comprehensive understanding of the research approach. Following this, the results will be presented and discussed in detail, offering insights into the acceptance or rejection of the hypotheses, thus contributing to a nuanced interpretation of the study's findings.

II. HYPOTHESES AND STUDY FRAMEWORK

A study conducted by [9] a thorough literature review to discern influential factors affecting SMEs' adoption of Cloud computing. From this examination, they devised a comprehensive framework, identifying 17 critical hypotheses encapsulating the most pressing challenges. Consequently, they proposed the 'Adoption of Cloud Computing Model by Saudi Arabian SMEs' (ACCM-SME), illustrated in Figure 1. Moreover, the framework was designed to represent the challenges in four primary contexts to investigate the fundamental difficulties in Cloud computing adoption. These contexts are: Technological, Organisational, Environmental and Social where each has a set of associated hypothesis linking factors to intention to adopt Cloud computing. These contexts and hypothesis are discussed in [9] but repeated here for convenience.

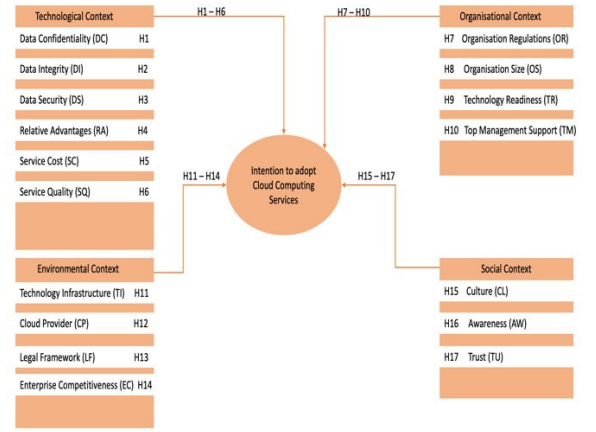


Fig. 1. The study framework (ACCM-SME)[9].

A. Technological context

The technology context is concerned with investigating the **technical factors** that influence an SME's decision to adopt Cloud computing services. Therefore, the following six factors (*Data Confidentiality (DC)*, *Data Integrity (DI)*, *Data Security (DS)*, *Relative Advantages (RA)*, *Service Cost (SC)*, and *Service Quality (SQ)*) investigate the most actual elements that impact the adoption process of Cloud computing services in Saudi SMEs. Then, based on the that six hypotheses have been developed as shown in the Table 1.

TABLE 1. TECHNOLOGICAL CONTEXT HYPOTHESES [9].

H	Hypothesis Statement
H1 → DC	“Increasing the data confidentiality of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H2 → DI	“Increasing the data integrity of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H3 → DS	“Increasing the data security of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H4 → RA	“The perceived relative advantages of Cloud computing have a positive effect that increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H5 → SC	“Decreasing the cost of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H6 → SQ	“Increasing the services quality of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”

B. Organisational context

The organisational context is concerned with investigating the **organisational factors** that influence an SME's decision to adopt Cloud computing services. Therefore, the following four factors (*Organisation Regulations (OR)*, *Organisation Size (OS)*, *Technology Readiness (TR)*, and *Top Management support (TM)*) relate to corporate practices that are built on the foundations and concepts of SMEs in Saudi Arabia to adopt Cloud services. Then, based on the that four hypotheses have been developed as shown in the Table 2.

TABLE 2. ORGANISATIONAL CONTEXT HYPOTHESES [9].

H	Hypothesis Statement
H7→ OR	“Increasing and updating the organisational regulations increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H8→ OS	“A smaller organisation size is more likely to increase the Saudi SMEs’ intention to adopt Cloud computing services ”
H9→ TR	“Increasing technology readiness increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H10→ TM	“Increasing the top management support of Cloud computing increases the Saudi SMEs’ intention to adopt Cloud computing services ”

C. Environmental context

The environmental context is concerned with investigating the **environmental factors** that influence an SME's decision to adopt Cloud computing. Therefore, the following four factors (*Technology Infrastructure (TI)*, *Cloud Provider (CP)*, *Legal Framework (LF)*, and *Enterprise Competitiveness (EC)*) discuss the most influential elements that may contribute to disrupting Saudi SMEs adopting Cloud computing services. Then, based on the that four hypotheses have been developed as shown in the Table 3.

TABLE 3. ENVIRONMENTAL CONTEXT HYPOTHESES [9].

H	Hypothesis Statement
H11→ TI	“Obtaining a high level of the technology infrastructure increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H12→ CP	“Increasing the number of Cloud providers within Saudi Arabia increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H13→ LF	“Obtaining an organised legal framework of Cloud computing increases Saudi SMEs’ intention to adopt Cloud computing services ”
H14→ EC	“Increasing the enterprise competitiveness increases the Saudi SMEs’ intention to adopt Cloud computing services ”

D. Social context

The Social context concerned with investigating the actual users’ **attitude and behaviour factors** that influence an SME's decision to adopt Cloud computing services. Therefore, the following three factors (*Culture (CL)*, *Awareness (AW)*, and *Trust (TU)*) discuss deeply all elements that influence the intention of Saudi SMEs toward adopting Cloud services from the perspective of the social context. Then, based on the that three hypotheses have been developed as shown in the Table 4.

TABLE 4. SOCIAL CONTEXT HYPOTHESES [9].

H	Hypothesis Statement
H15→ CL	“Increasing the technology culture for customers, IT managers, top managers, and employees increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H16→ AW	“Increasing the awareness of the customers, IT managers, top managers and employees increases the Saudi SMEs’ intention to adopt Cloud computing services ”
H17→ TU	“Increasing the trust of the customers, IT managers, top managers and employees increases the Saudi SMEs’ intention to adopt Cloud computing services ”

III. STUDY METHODOLOGY

A. Approach and Sample

This research employs a quantitative approach by conducting a comprehensive survey as the primary means of gathering initial data. The aim is to examine and assess the degree to which the adoption of Cloud computing services is effective among SMEs in Saudi Arabia. This study will discuss quantitative data analysis and provide a general overview of the fieldwork that was conducted to collect data from IT managers, top managers, and employees of SMEs in Riyadh, Saudi Arabia. An extensive examination of existing literature designed the survey questions. Then, three closed-ended questions were used as a minimum to measure each factor. This study used a 5-point Likert scale to measure participants' replies, with a rating of one indicating “strongly disagree” and five indicating “strongly agree”.

The target population selected in Riyadh exceeds 100,000 IT managers, top managers, and SMEs employees, according to a report from the Small and Medium Enterprise General Authority in the Kingdom of Saudi Arabia. The sample size can be calculated using various techniques following the determination of the confidence level beside the margin of error. Therefore, based on the [10] technique, the sample size for this study has been calculated; it is 384, with a level of confidence that attains 95% degree and a margin of error that attains 5% degree. The sample was chosen randomly in 2022 from among Saudi SMEs in Riyadh. The link to a web-based survey was made and sent to each SMEs sample member separately. Following completion, each response was saved immediately into the database.

B. Survey Implementation

The main goal of this part of the study is to examine the relationship of the developed hypotheses with the Saudi SMEs intention to adopt Cloud computing services. There was a total of 66 questionnaires used in this study; 11 were used to collect demographic data on the participants, and their respective companies, and the remaining 55 were used to assess the influence of various factors on the adoption of Cloud computing services among Saudi SMEs.

The questions and instructions that were asked in the primary survey were tested in a pilot study first, carried out in Saudi Arabia with the participation of IT managers, top managers, and staff members of Saudi SMEs. The pilot study was sent to the participants with the intention of gathering 20 responses. However, 27 responses were actually filled out and returned, which helped to evaluate the questionnaire draught. The primary objective of the pilot study was to determine whether the questionnaires and instructions could be easily understood and filled out without any problems and whether the allotted time was sufficient to complete the questionnaires. Therefore, the pilot study's findings were beneficial because they revealed great comments from the participants but did not result in any change in the questions.

Therefore, the main study was conducted and distributed from the 5th of September 2022 to the 3rd of December 2022 to the IT managers, top managers, and SMEs employees in Riyadh City, Saudi Arabia. Over 500 questionnaires were randomly distributed over three months; 440 (88%) out of

500 were filled out and returned. However, of these responses returned, 28 (6,36%) were considered invalid data as they had missing data for variables and would affect the analysis process, and 412 (93.64%) were complete and accurate, so it is used.

IV. STUDY RESULTS

A. Confirmatory Factor Analysis (CFA)

Confirmatory Factor Analysis is considered an effective pattern that verifies the factors through variable load testing. It is carried out and emphasized to determine the relationship between observed variables and the theoretical definition of the factors.

Before receiving any results, the researcher must identify the correct factor to decide whether or not to accept the theory being investigated by picking the proper theory being researched. Therefore, following the CFA, the results influence the decision on the theory, and the procedure logically has two options: supported theory (accept) or unsupported theory (reject) [11].

Furthermore, when the researcher decides to employ the CFA, it should evaluate each construct measurement model's unidimensionality, reliability, and validity. To determine if each variable measures a single underlying feature, a latent measurement model for each construct, including external and internal factors, should be established. This is a critical stage in carrying out a technique to determine unidimensionality. It is essential to establish unidimensionality before proceeding with structural equation modelling, which entails removing any signs with factor loadings less than 0.60 [12].

B. Construct Validity and Reliability

The assessments used to evaluate SMEs' adoption of Cloud computing services were developed after thorough literature research and then adjusted to this study. The main purpose of any assessment tools employed is to ensure that the study model is very accurate and has reliable data that can fit with it.

Three types of validity—convergent, discriminant, and reliability—were utilised to check the study model. Moreover, the measurement model's Composite Reliability (CR) and Cronbach's alpha were used to test this study's construct validity and reliability. CR and Cronbach's alpha should exceed the threshold value, which is 0.70 [13]. Table 5 shows that the construct validity and reliability results met the required level as Composite Reliability (CR) and Cronbach's alpha were over 0.70.

Convergent validity evaluation means the model has to have a good fit, and the AVE value needs to be more than 0.5. [14]. A high AVE indicates a robust relationship between latent and single variables. According to [15], [16], they proposed values for most indices as thresholds for them, which are as follows: CMIN/DF value should be between 1 and 5, CFI value should be ≥ 90 , RMSEA should be ≤ 0.06 , and SRMR should be ≤ 0.08 , and when the thresholds are met, the model fit is considered acceptable.

In addition, according to [17], in order to determine whether or not a model is a good fit for the data that was obtained from more than 250 participants, the chi-square statis-

tic must be less than 3. In this particular investigation, information was gathered from a total of 412 participants, resulting in a chi-square value of 1764.4. All of the model fit measurements have reached the excellent values, including CFI = 0.956, SRMR = 0.053, RMSEA = 0.031. This means that the needed degree of model fit at this stage has been obtained, as shown in Table 5.

TABLE 5: CONSTRUCT VALIDITY AND RELIABILITY.

Factors	Cronbach (≥ 0.7)	CR	AVE	MSV
Data Confidentiality (DC)	0.808	0.811	0.591	0.157
Data Integrity (DI)	0.748	0.749	0.500	0.146
Data Security (DS)	0.770	0.778	0.539	1.263
Relative advantages (RA)	0.747	0.749	0.500	0.456
Service Cost (SC)	0.783	0.784	0.547	0.804
Service Quality (SQ)	0.787	0.787	0.552	0.410
Organisation Regulation (OR)	0.821	0.820	0.603	0.223
Organisation Size (OS)	0.807	0.807	0.583	1.058
Technology Readiness (TR)	0.779	0.777	0.539	0.580
Top Management Support (TM)	0.778	0.778	0.540	0.014
Technology Infrastructure (TI)	0.781	0.784	0.548	0.372
Cloud provider (CP)	0.818	0.814	0.595	0.470
Legal Framework (LF)	0.787	0.808	0.587	0.275
Enterprise Competitiveness (EC)	0.788	0.787	0.554	0.145
Culture (CL)	0.771	0.770	0.529	0.281
Awareness (AW)	0.751	0.751	0.504	0.219
Trust (TU)	0.810	0.815	0.597	0.149
Intention to Adopt Cloud computing services (ITAC)	0.859	0.843	0.571	0.736

When describing the ability to differentiate one construct from others of the same kind, the term "discriminant validity" is used. This study assesses the discriminant validity using the [14] recommendation to link the value of squared roots estimates with Average Variance Extracted (AVE) values. According to [17] guidelines, the evaluation of discriminant validity comprises three distinct benchmarks, namely the Maximum Shared Squared Variance (MSV), the Fornell-Larcker test, and the Average Shared Squared Variance (ASV). Table 6 illustrates that the obtained square root of the constructs' Average Variance Extracted (AVE) value was greater than the values below the diagonal estimates.

The current stage endeavours to conduct a comparative evaluation of the value of absolute relationship among constructs and the square root of the average variance extracted

TABLE 6: DISCRIMINANT VALIDITY

	DC	DI	DS	RA	SC	SQ	OR	OS	TR	TM	TI	CP	LF	EC	CL	AW	TU	ITAC
DC	.768																	
DI	.241	.707																
DS	.021	.016	.734															
RA	.227	.214	.265	.707														
SC	.247	.202	.374	.494	.740													
SQ	.180	.191	.206	.470	.450	.743												
OR	.143	.167	.273	.377	.348	.308	.777											
OS	.225	.212	.171	.473	.432	.383	.381	.763										
TR	.249	.211	.370	.573	.725	.462	.375	.455	.734									
TM	-.101	-.088	-.111	.000	-.103	-.109	-.055	-.026	-.122	.734								
TI	.272	.226	.190	.486	.465	.352	.298	.395	.460	-.113	.740							
CP	.249	.229	.411	.601	.610	.519	.431	.487	.620	-.110	.440	.771						
LF	.275	.217	.287	.513	.589	.502	.350	.470	.603	-.093	.486	.593	.766					
EC	.017	.034	.828	.216	.300	.198	.217	.162	.302	-.099	.162	.344	.270	.744				
CL	.179	.186	.206	.518	.424	.428	.417	1.02	.483	.018	.376	.502	.464	.174	.727			
AW	.227	.191	.154	.402	.506	.326	.263	.350	.462	-.124	.395	.440	.365	.137	.375	.709		
TU	.154	.145	.215	.358	.405	.257	.308	.371	.393	-.034	.284	.469	.429	.188	.363	.278	.772	
ITAC	.353	.323	.287	.715	.639	.585	.537	.655	.698	-.037	.590	.734	.687	.242	.679	.520	.515	.758

(AVE) for each construct. According to the criterion established by [18], if the relationship are less than the square root of the construct's average variance extracted, then it is improper to argue that a model has attained discriminant validity. However, the model demonstrates discriminant validity as evidenced by the "diagonal cells" of the Average Variance Extracted's square root exhibiting a greater value than the values below the diagonal estimates among constructs. Therefore, in the next section, this paper will discuss the data analysis intensely by showing the descriptive data first, then determining which hypothesis has been accepted and which was rejected through structural equation modelling testing.

V. DATA ANALYSIS

It is essential to ensure the validity and reliability of the model through many measurement processes to identify its quality before going to further statistical analysis. Therefore, this section shows the details about the instruments that have been used in the test of the Confirmatory Factor Analysis (CFA). For the consistent validity test, the Statistical Software for the Social Sciences (SPSS) version 28 was employed. At the same time, the CFA assessed the model's fitness by employing the software of Analysis of Moment Structure (AMOS) version 28. However, before going to test hypotheses, the use of these instruments assists in giving proof of the validity and reliability of the model's unidimensional, discriminant, and convergent properties. The CFA findings indicated 18 latent constructs, which are the intention to Adopt Cloud computing services (ITAC), Data Confidentiality (DC), Data Integrity (DI), Data Security (DS), Rel-

ative advantages (RA), Service Cost (SC), Service Quality(SQ), Organisation Regulation (OR), Organisation Size(OS), Technology Readiness (TR), Top Management Support (TM), Technology Infrastructure (TI), Cloud provider (CP), Legal Framework (LF), Enterprise Competitiveness (EC), Culture (CL), Awareness (AW), Trust (TU). To investigate the perceived efficacy of using Cloud computing services by Saudi SMEs, 18 factors and 55 items were used. Furthermore, the loading factors for the majority of individual items were larger than 0.60. To ensure that a measurement model procedure is unidimensional, any items to be employed must have a loading factor of 0.50 or above, while any things with a loading factor of 0.50 or below must be disregarded and removed from consideration [14], [19]. As a result, after conducting an analysis of data and developing the model by AMOS software, no items were holding loading under 0.50, which meets the threshold of most references.

A. Descriptive statistics

This study has collected 440 responses. However, of these responses collected, 28 (6.36%) were considered invalid data as they had missing data for variables and would affect the analysis process, and 412 (93.64%) were complete and accurate.

After analysing the data on participants' ages, the result indicated that 15.5% are between 18-25 years, 36.9% are between 26-35 years, 31.1% are between 36-45, and 16.5% are equal or greater than 46 years. The result of this study found that 23.1% of participants are female and 76.9% are male. The result of the data regarding the levels of Education of the participants revealed that the majority of them, which is 195

(47.3%) participants, have a bachelor's degree, 107 (26%) participants have a diploma degree, and 72 (26%) participants have a master's degree, and 30 (7.3%) participants have a high school or less, and 8 (1.9%) participants have a Ph.D. degree.

Regarding the participants' positions, the data revealed that 82 (19.9%) of the directors of organizations participated in this study, along with 22 (5.3%) IT directors, 48(11.7%) IT staff, 220 (53.4%) employees, and 40 (9.7%) selected the others. The number of years of experience for participants was 29 (7.1%) with one year of experience or less, 160 (38.8%) with one year of experience to three years, 90 (21.8%) with four years of experience to five years, and 133 (32.3%) with more than six years of experience.

Regarding the distribution of SMEs participants in this study on the different sizes of Saudi SMEs, the data revealed 228 (55.3%) participants from medium-sized enterprises, 136 (33%) participants from small-sized enterprises, and 48 (11.7%) participants from micro-sized enterprises (see Table 7).

The business sectors of participants who participated in this study, the data showed that the largest number of participants in the study worked in Information Technology and Communication Sector 48 (11.7%), followed by the Retail Sector 42 (10.2%), then Education and Training Sector 39 (9.5%) and Financial Sector 39 (9.5%). Then, the Construction and Contracting Sector 36 (8.7%), the Administrative Services Sector 28 (6.8%), the Energy Sector 27 (6.6%), the Health Sector 25 (6.1%), the Manufacturing Sector 23 (5.6%), Estate and Utilities Sector 23 (5.6%), Transportation and Storage Sector 22 (5.3%), Others 20 (4.9%) which is a type of business that is not categorized under a specific sector or its sector is not listed in this study, Food Sector 15 (3.6%), Home Services Sector 14 (3.4%), and Sports Sector 11 (2.7%) (see figure 2).

The participants were asked whether their organizations had adopted the Cloud services. 92 (22.3%) answered yes, and 320 (77.7%) responded no. However, those 92 participants who answered yes to their organizations adopting Cloud computing services were asked the additional question of what type of Cloud their organization used. Then out of 92, 20 (21.3%) used public Cloud, 27 (28.7%) used private Cloud, 1 (1.1 %) used community Cloud, 6 (6.4%) used hybrid Cloud, and 38 (41.30%) they do not know which the Cloud is their organization used.

Overall, the descriptive data of the participants in this study shows that data were collected from various participants and various types of Saudi SMEs, as Figure 2 shows.

TABLE7: DEMOGRAPHICS DATA.

Item	Type	Frequency	Percent
Collected Data	Valid	412	93.64
	Missing	28	6.36
Age	18-25	64	15.5
	26-35	152	36.9
	36-45	128	31.1
	>=46	68	16.5
Gender	Female	95	23.1
	Male	317	76.9
Education	High School or less	30	7.3
	Diploma	107	26
	Bachelor	195	47.3
	Master	72	17.5
	Ph.D.	8	1.9
Job positions	Director of organisation	82	19.9
	IT director	22	5.3
	IT staff	48	11.7
	employees	220	53.4
	others	40	9.7
Years of experience	1 year or less	29	7.1
	1-3 years	160	38.8
	4-6 years	90	21.8
	More than 6 years	133	32.3
SMEs sizes	micro-sized	48	11.7
	small-sized	136	33
	medium-sized	228	55.3




Figure 2: The participants' sectors.

TABLE 8: MODEL FIT MEASURES.

Measure	Chi-square (CMIN) / Degrees of freedom (DF)	Comparative fit index (CFI)	Standardised root mean square residual (SRMR)	Root mean square error of approximation (RMSEA)	P of Close Fit (PClose)
Estimate	1.384	0.956	0.053	0.031	1.000
Threshold	Between 1 and 5	> 0.95	< 0.08	< 0.06	> 0.05
Interpretation	Good fit	Good fit	Good fit	Good fit	Good fit

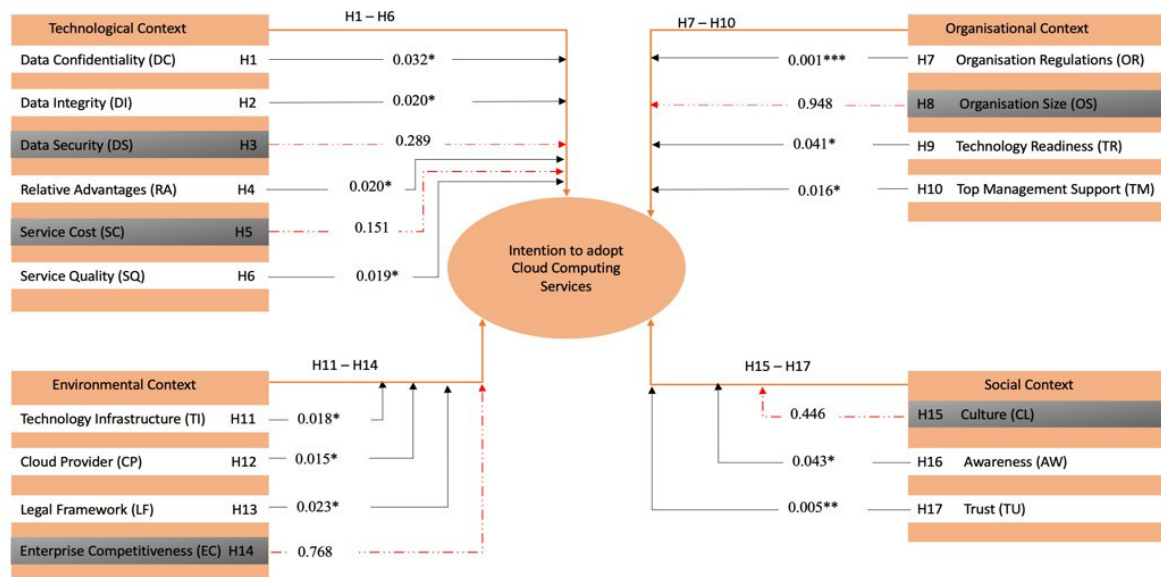


Fig. 3. The study framework (ACCM-SME) after analysis.

B. Structural Equation Model (SEM)

The structural equation modelling testing shows the relationship between empirically tested constructs and whether it is significantly related or not to each other. The structural equation model and measurement models are not identical; instead, there is a distinction between the two. In the measurement model, the focus is placed on assessment, which demonstrates the links between constructs. Still, in the structural model, measured variables are used to illustrate how significant the relationships between constructs are. According to [20], the hypothesis testing object through SEM is to identify the extent to which predictors, also known as independent variables, contribute to explaining the dependent variables being studied. Overall, the model identified the following exogenous (independent) constructs: Data Confidentiality (DC), Data Integrity (DI), Data Security (DS), Relative advantages (RA), Service Cost (SC), Service Quality (SQ), Organisation Regulation (OR), Organisation Size (OS), Technology Infrastructure (TI), Cloud provider (CP), Legal Framework (LF), Enterprise Competitiveness (EC), Culture (CL), Awareness (AW), and Trust (TU). The endogenous (dependent) construct is then identified as the Intention to Adopt Cloud computing services (ITAC). Therefore, the model indices in this stage assure that it is at a good level of fitting as follows: ($\chi^2 = 1764.4$, $df = 1275$, $CFI = 0.956$, $SRMR = 0.053$, and $RMSEA = 0.031$) as shown in the table 8.

The researchers used SEM with AMOS 28 to build the model and check whether it is good fitting. Therefore, the path analysis was assessed based on the structural model fit criteria. This research developed the framework (ACCM-SME), which consisted of 17 different relationships that were hypothesised based on the literature. The empirical data analyses revealed that 12 out of the 17 hypotheses tested had significant and positive influence, and five did not satisfy the

TABLE 9: PATH ANALYSIS FOR ALL CONSTRUCTS.

NOTE: (“***=SIGNIFICANCE AT THE 0.001 LEVEL, **=SIGNIFICANCE AT THE 0.01 LEVEL AND *=SIGNIFICANCE AT THE 0.05 LEVEL”)

Structural Relation	Regression Weight	Standard Error (S.E.)	Critical ratio (C.R.)	P value	Result
ITAC ← DC	.094	.044	2.149	.032*	Accepted
ITAC ← DI	.125	.054	2.326	.020*	Accepted
ITAC ← DS	-.061	.057	-1.061	.289	Rejected
ITAC ← RA	.136	.058	2.323	.020*	Accepted
ITAC ← SC	-.191	.134	-1.428	.151	Rejected
ITAC ← SQ	.116	.050	2.344	.019*	Accepted
ITAC ← OR	.135	.035	3.855	.001***	Accepted
ITAC ← OS	-.013	.202	-.065	.948	Rejected
ITAC ← TR	.295	.144	2.044	.041*	Accepted
ITAC ← TM	.066	.027	2.411	.016*	Accepted
ITAC ← TI	.115	.048	2.364	.018*	Accepted
ITAC ← CP	.131	.054	2.440	.015*	Accepted
ITAC ← LF	.154	.068	2.273	.023*	Accepted
ITAC ← EC	.020	.068	.295	.768	Rejected
ITAC ← CL	.145	.190	.762	.446	Rejected
ITAC ← AW	.139	.068	2.028	.043*	Accepted
ITAC ← TU	.105	.037	2.801	.005**	Accepted

criteria specified for this research, which have been rejected, as shown in table 9, and in updated framework see figure 3.

Most factors of this empirical study affect the intention of Saudi SMEs to adopt Cloud computing services except five of them, which are Data Security (DS) (P=0.289), Service Cost (SC) (P=0.151), Organisation Size (OS) (P=0.948), Enterprise Competitiveness (EC) (P=0.768), and Culture (CL) (P=0.446) (see table 9). This will be discussed in more detail in the following section.

VI. DISCUSSION

This study has developed a framework in the context of adopting Cloud computing services by Saudi SMEs. Then, it was tested through empirical data to identify the relationship through factors hypothesis based on four contexts, which are: Technological context, Organisational context, Environmental context, and social context.

A. Technological context

The technological context is an essential element of economic enhancement, and it plays an indispensable role in transforming the organisations' traditional work into more reliable and advanced technological pathways. This requires a contemporary and robust technology infrastructure.

The IT background and usage of Cloud computing in Saudi SMEs have been addressed by recent studies [21], [22], [23]. However, confidentiality, privacy, data control, the service cost, a lack of Cloud knowledge, and the integrity of the data have not been adequately considered in the recent studies that have been conducted in the area of data security [6], [24].

To fill the knowledge gap mentioned by the previous authors in their recommendations for further investigation, this study investigated to identify all of the significant factors that influence Cloud computing adoption in Saudi SMEs in the context of technology. Then, the result of our study revealed that data confidentiality and data Integrity are considered important factors that are significantly related to the intention to adopt Cloud computing by Saudi SMEs, and that results were consistent with the results of [25].

Data security as a significant term was not a concern for Saudi SMEs, which may lay for some reasons, one of them being a recent requirement that the Saudi government forces all businesses to store their data locally in Saudi.

Service quality and relative advantage were found to play a role in our investigation. They positively affect the adoption of Cloud computing services, and [26] showed the significance of service quality and relative advantage along with their benefits toward Cloud adoption.

Service cost found to be not posed a concern or a challenge. It was expected that lower service costs would encourage Cloud computing use; this has not been the case, and it may be due to the concern of Saudi SMEs about quality, whatever the cost is.

B. Organisational context

The organisational context investigates four factors revealed after a deep dive in the literature: Organisation regulation, Organisation size, Technology readiness, and Top management support. The organisational context is an essential and indispensable part that plays a significant role in adopting Cloud computing and impacts most developing countries, including Saudi Arabia [27].

Our results revealed that organisation regulation is a crucial factor and has a relationship with the intention of adopting the Cloud and positively affected Saudi SMEs, consistent with [28], who confirmed that keeping updated organisation regulations reinforces the adoption of the Cloud.

In addition, organisation size was found not to be a relationship with Saudi SMEs' intention to adopt Cloud comput-

ing services, and this result refuted what has been found in this study regarding the importance of organisation size in the adoption of Cloud [29].

From the perspective of technology readiness, it positively affects the adoption of the Cloud. The path coefficient is relative, suggesting that this element has a noteworthy impact on Saudi SMEs' decisions to move their operations to the Cloud, which is consistent with the result [30].

Top management support is considered an essential factor related to the decision that may be made regarding Cloud adoption. Other studies have investigated how top management affects the adoption of Cloud computing services and discovered that they relate to each other positively [31]. Therefore, this study's results found that the top management support factor significantly positively influences Saudi SMEs' adoption of Cloud computing.

C. Environmental context

The environmental context is considered a key that leads organisations to adopt Cloud computing services as it represents significant factors that are part of the life cycle of organisations. Technology infrastructure is one of them that has been investigated in this study. A study has assessed how the private Cloud with solid technology infrastructure helps organisations migrate from traditional technology to the Cloud. It found that it is beneficial and helps the organisation carry many virtual services, houses more than 35 services without any disruption, and is capable of more [32]. Our results were consistent with that and confirmed the technology infrastructure substantially impacts the Saudi SMEs' intention to adopt Cloud computing services as the critical ratio demonstrated a significant connection between them.

A Cloud provider is the second factor in the environmental context that seems to impact SMEs' adoption of Cloud services due to Saudi government restrictions, which prohibit dealing with any Cloud provider whose data centres are outside of Saudi, which resulted in a limited number of Cloud providers within Saudi. An increasing number of Cloud providers within Saudi increasing SMEs' intention to adopt the Cloud is the hypothesis that was made regarding the Cloud provider factor. It was correct, and it was found that a positive relationship exists between the Cloud provider and the adoption of Cloud computing in Saudi SMEs. [33] confirmed that organisations in Saudi Arabia trust the Cloud provider based in Saudi Arabia, which increases their confidence and results in adopting the Cloud.

Regarding the third factor in the environmental context of the legal framework, [34] found that when the legal framework and strategies within organisations are revised continuously, it leads to getting good quality and pushing for adopting Cloud services. Therefore, this study found a positive link between the Legal framework and the adoption of Cloud computing in Saudi SMEs.

Then, competitive pressure, the fourth factor in the environmental context, was investigated to check its effects. According to [35], competitive pressure is an influencer factor in adopting Cloud. Moreover, our findings were against that and showed no concern for enterprise competitiveness, which resulted in no impact on the intention of Saudi SMEs to adopt Cloud services.

D. Social context

The social context plays a significant role in the success of adopting new technologies, especially the adoption of Cloud computing services. Social context is one of the most essential parts of understanding any problem, and the most critical factors have been found, including societal behaviour (culture) and awareness [36]. However, this part of the study investigates three social context factors identified after a deep search in the literature about the factors that may affect the decision to adopt Cloud computing services: culture, awareness, and trust.

This study showed no relationship or influence between culture and intention to adopt Cloud computing services by Saudi SMEs as follows ($\beta = .145$, $p < 0.446$). This refuted the result of [37], who found a negative relationship and indirect impact of cultural factors on the decision of Saudi SMEs to adopt the Cloud, and [38], who found that culture significantly impacts the Cloud computing adoption decision in developing countries.

In addition, the result of this study regarding awareness revealed that the route from awareness to intention to adopt Cloud computing services had a positive effect ($\beta = .139$, $p < 0.043$), indicating that all derived values, including the crucial ratio for each item, were substantially within the range. This implies that Saudi SMEs feel embracing Cloud computing depends on their employees' awareness. Saudi SMEs should be aware of adopting Cloud services effectively to fulfil organisations' demands. Consequently, the data analysis results of the awareness component showed a positive relationship with Cloud computing adoption. Indeed, the study results are consistent with the outcome of [36], which confirmed that people's awareness should be considered as it is regarded as a key that helps accelerate organisations' transition toward Cloud adoption.

Regarding the trust factor, the study results found a significant positive relationship between trust and the adoption of Cloud computing in Saudi SMEs. Thus, the noteworthy, related relationship, as shown by the critical value obtained between trust and intention to adopt Cloud computing services ($\beta = .105$, $p < 0.005$). This indicates the trust factor significantly impacts the level to which Saudi SMEs adopt Cloud computing. Many authors agree that trust is essential to maintain transparency between organisations and their employees, pushing them to adopt Cloud computing services, for example [27], [39].

Therefore, Data Security (DS) ($P=0.289$), Service Cost (SC) ($P=0.151$), Organisation Size (OS) ($P=0.948$), Enterprise Competitiveness (EC) ($P=0.768$), and Culture (CL) ($P=0.446$) were the independent factors that did not significantly affect Saudi SMEs' adoption of Cloud computing, as shown in Table 9.

The null hypothesis H3, "Increasing the data security of Cloud computing increases the Saudi SMEs' intention to adopt Cloud computing services", was rejected with a value of ($P=0.289$), indicating that the data security had no impact on the decision of the Saudi SMEs to adopt Cloud computing services. Similarly, the null hypothesis H5, "Decreasing the cost of Cloud computing increases the Saudi SMEs' intention to adopt Cloud computing services", was rejected with a value of ($P=0.151$), indicating that the cost associated with

Cloud computing had no significant impact on the decisions of Saudi SMEs to adopt Cloud computing.

This outcome may be attributed to the Saudi government's assistance and expertise in information technology (IT), which has helped the private sector improve its data security practices. In addition, the fact that Saudi organisations do not care about saving money when it comes to transforming into new technologies is evidence that they value innovation and competitiveness above all else.

Moreover, the null hypothesis H8, "A smaller organisation size is more likely to increase the Saudi SMEs' intention to adopt Cloud computing services", was rejected with the value of ($P=0.948$), showing that the organisation size does not influence the intention of Saudi SMEs toward adoption of Cloud computing services. In addition, the value of ($P=0.768$) proving that the null hypothesis H14, "Increasing the enterprise competitiveness increases the Saudi SMEs' intention to adopt Cloud computing services", was rejected, indicating that the enterprise competitiveness does not have an impact on the intention to adopt Cloud computing services by Saudi SMEs. The culture factor seems does not affect the intention of Saudi SMEs to adopt Cloud computing services since the null hypothesis H15, "Increasing the technology culture for customers, IT managers, top managers, and employees increases the Saudi SMEs' intention to adopt Cloud computing services", was rejected with the value ($P=0.446$).

This empirical data adds to our knowledge of the essential elements from the viewpoint of SMEs, for whom Cloud computing adoption become a big concern. This research addresses a knowledge gap and expands SMEs' grasp of Cloud computing in Saudi Arabia. The difference that may be noticed between the result of this study and other studies is that variations between nations and the populations studied might account for discrepancies with previous research.

VII. CONCLUSIONS

Nowadays, Cloud computing has become an essential technology service on a par with the service of electricity, gas, and water. It plays a significant role in making the process and storage of data very flexible and on everyone's hands as it offers pay-per-use service. The result of the study may contribute to helping SMEs in the future by making them conscious of the challenges that face Cloud adoption in the early stage.

Further, the findings of this study may contribute to assisting organisations and their leadership to understand the accurate perception of those organisations. In addition, it may help SME reformulate their thinking and strategies, around Cloud adoption, which is consistent with factors that this study has proved as challenges SMEs face.

Due to the technological advancements facilitated by Cloud computing, such as improved infrastructure flexibility and efficiency, enhanced data processing capabilities, and heightened automation, alongside the associated benefits of cost reduction and operational optimization, many organizations are keen to implement Cloud solutions. This study empirically explores the factors influencing organizations' adoption of Cloud services, shedding light on both the challenges and obstacles encountered across technological, organisational, environmental, and social contexts.

Specifically, this study investigates seventeen factors (Data Confidentiality, Data Integrity, Data Security, Relative Advantages, Service Cost, Service Quality, Organisation Regulations, Organisation Size, Technology Readiness, Top Management support, Technology Infrastructure, Cloud Provider, Legal Framework, Enterprise Competitiveness, Culture, Awareness, and Trust) affecting SMEs' Cloud computing adoption within these four contexts. First, a systematic literature review has been conducted [9], to find the knowledge gap around factors that might influence adoption of Cloud computing in the contexts mentioned. Critical factors that pose challenges and obstacles to Saudi SMEs have thus been identified and modelled in the ACCM-SME framework, Figure 1. Then, the research framework was developed based on the TOE model blended with DOI theory in different contexts (technological, organisational, environmental, and social). Based on that, 17 hypotheses (covering all factors) were constructed leading to a compressive survey development with 66 items; 11 querying demographic data, while 55 examined how the factors influence Cloud computing adoption empirically.

The study survey was distributed to over 500 participants in Saudi SMEs in Riyadh city, including Top management, IT directors, and employees. Four hundred and twelve participants have fully answered the survey. As a result, after analysing the data, we found that the factors: Data Security (DS), Service Cost (SC), Organisation Size (OS), Enterprise Competitiveness (EC), and Culture (CL) showed no impact on the Saudi SMEs to adopt Cloud computing services. In contrast, the rest of the identified factors impacted Cloud adoption.

This research shows that through the ACCM-SME framework (based on the merging of two established theories called the TOE model and DOI theory), we were able to examine SMEs' use of Cloud computing from a more comprehensive technological, organisational, environmental, and social vantage point, which may give helpful insights for businesses.

The research also clarified and highlighted the importance of leaders in fostering an environment conducive to developing systems that need the involvement of organisational leaders with the skills to overcome obstacles in adopting and implementing suitable information technology. Several factors that might encourage and reinforce businesses to switch to Cloud-based systems have been explored in this research.

Through the ACCM-SME framework we were able to establish both factors that had no influence and those that influenced adoption of Cloud computing from an SME perspective (Table 9, Figure 3). This gave us useful insights as to why SMEs are not yet moving to the Cloud. These are crucial findings for SMEs, particularly in Saudi Arabia, where a better understanding of relevant factors may lead to an increase in SME intention to adopt Cloud computing. As such, the ACCM-SME framework will serve as a useful resource for SMEs when they intend to start adopting Cloud services.

A. Study limitations

This study's results indicate how SMEs in future can both benefit from and adopt Cloud computing. Nonetheless, various limitations have been highlighted, giving an opportunity

to investigate them. Directors, IT directors, and workers of SMEs were among the survey respondents for this research. As a result, there are two limitations that were out of the study scope. First, the study has not considered the views of consumers or Cloud computing service providers. Secondly, it does not consider job opportunities that may arise from the adoption of Cloud computing by the SME, and which may influence their choice to adopt. These limitations may be overcome in the future as new variables of study.

B. Future work

In the upcoming stage of our research, these results will also be validated by qualitative data. This study will be extended to the other factors that influence the process of Cloud adoption, such as the perception of Cloud providers and job opportunities that Cloud adoption may disparage. It is worth noting that the challenges that face SMEs to adopt Cloud computing services may differ from one country to another due to the different environments and the styles of organisations. Finally, to increase the efficiency of future studies, the researchers should think about acquiring samples with various viewpoints of all partners who are related to SMEs.

REFERENCES

- [1] F. Dahlqvist, M. Patel, A. Rajko, and J. Shulman, "Growing opportunities in the Internet of Things," *McKinsey*, July, 2019.
- [2] M. Attaran, "Cloud computing technology: leveraging the power of the internet to improve business performance," *Journal of International Technology and Information Management*, vol. 26, no. 1, pp. 112–137, 2017. DOI: <https://doi.org/10.58729/1941-6679.1283>.
- [3] R. Almajalid, "A survey on the adoption of cloud computing in education sector," *arXiv preprint arXiv:1706.01136*, 2017. <https://doi.org/10.48550/arXiv.1706.01136>.
- [4] W. Klug and X. Bai, "Factors affecting cloud computing adoption among universities and colleges in the United States and Canada," *Issues in Information Systems*, vol. 16, no. 3, 2015.
- [5] R. Sahandi, A. Alkhalil, and J. Opara-Martins, "Cloud computing from SMEs perspective: a survey based investigation," *Journal of Information Technology Management*, vol. 24, no. 1, pp. 1–12, 2013.
- [6] M. O. Alassafi, A. Alharthi, R. J. Walters, and G. B. Wills, "Security risk factors that influence cloud computing adoption in Saudi Arabia government agencies," in *2016 International Conference on Information Society (i-Society)*, IEEE, 2016, pp. 28–31. DOI: [10.1016/j.tele.2017.04.010](https://doi.org/10.1016/j.tele.2017.04.010).
- [7] B. S. Alghamdi, M. Elnamaky, M. A. Arafah, M. Alsabaan, and S. H. Bakry, "A Context Establishment Framework for Cloud Computing Information Security Risk Management Based on the STOPE View," *IJ Network Security*, vol. 21, no. 1, pp. 166–176, 2019. DOI: [10.6633/IJNS.201901.21\(1\).21](https://doi.org/10.6633/IJNS.201901.21(1).21).
- [8] Trade, "Saudi Arabia - Information and Communications Technology," International Trade Administration. Accessed: Nov. 09, 2021. [Online]. Available: <https://www.trade.gov/country-commercial-guides/saudi-arabia-information-and-communications-technology>
- [9] M. Alqahtani, N. Beloff, and M. White, "A new adoption of cloud computing model for Saudi Arabian SMEs (ACCM-SME)," in *Proceedings of SAI Intelligent Systems Conference*, Springer, 2022, pp. 192–210.
- [10] M. Saunders, P. Lewis, and A. Thornhill, *Research methods for business students*. Pearson education, 2009.
- [11] J. L. Perry, A. R. Nicholls, P. J. Clough, and L. Crust, "Assessing model fit: Caveats and recommendations for confirmatory factor analysis and exploratory structural equation modeling," *Meas Phys Educ Exerc Sci*, vol. 19, no. 1, pp. 12–21, 2015. <https://doi.org/10.1080/1091367X.2014.952370>.
- [12] W. M. A. B. W. Afthanorhan, S. Ahmad, and I. Mamat, "Pooled Confirmatory Factor Analysis (PCFA) using structural equation modeling on volunteerism program: A step by step approach," *International Journal of Asian Social Science*, vol. 4, no. 5, pp. 642–653, 2014.
- [13] T. Schmiedel, J. Vom Brocke, and J. Recker, "Development and validation of an instrument to measure organizational cultures' support of

- Business Process Management,” *Information & Management*, vol. 51, no. 1, pp. 43–56, 2014. <https://doi.org/10.1016/j.im.2013.08.005>.
- [14] J. F. Hair, W. C. Black, B. J. Babin, R. E. Anderson, and R. L. Tatham, “Multivariate data analysis (Vol. 6),” 2006.
- [15] R. Kline, “Exploratory and confirmatory factor analysis,” in *Applied quantitative analysis in education and the social sciences*, Routledge, 2013, pp. 183–217.
- [16] Z. Awang, A. Afthanorhan, and M. A. M. Asri, “Parametric and non parametric approach in structural equation modeling (SEM): The application of bootstrapping,” *Mod Appl Sci*, vol. 9, no. 9, p. 58, 2015. DOI:10.5539/mas.v9n9p58.
- [17] J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, *Multivariate data analysis: Pearson new international edition PDF eBook*. Pearson Higher Ed, 2013.
- [18] C. Fornell and D. F. Larcker, “Evaluating structural equation models with unobservable variables and measurement error,” *Journal of marketing research*, vol. 18, no. 1, pp. 39–50, 1981. <https://doi.org/10.2307/3151312>.
- [19] E. C. Papanastasiou, “Factor structure of the ‘Attitudes Toward Research’ scale,” *Statistics education research journal*, vol. 4, no. 1, pp. 16–26, 2005. DOI:10.1037/t64085-000.
- [20] A. F. Hayes and K. J. Preacher, “Statistical mediation analysis with a multicategorical independent variable,” *British journal of mathematical and statistical psychology*, vol. 67, no. 3, pp. 451–470, 2014. <https://doi.org/10.1111/bmsp.12028>.
- [21] D. Alrubaishi and P. Robson, “Innovation in Saudi Family SMEs: The Role of Social Capital and Family Involvement,” *International Review of Entrepreneurship*, vol. 17, no. 1, 2019.
- [22] M. R. Faridi and A. Malik, “Customer engagement technology in sme’s in Saudi Arabia: Does it ensue in disturbance or disruption,” *International Journal of Entrepreneurship*, vol. 23, no. 1, pp. 1–8, 2019.
- [23] K. Mezghani and M. A. Almansour, “Study of Intentions to Use Cloud CRM Within Saudi SMEs: Integrating TAM and TPB Frameworks,” in *Business Transformations in the Era of Digitalization*, IGI Global, 2019, pp. 33–50. DOI: 10.4018/978-1-5225-7262-6.ch003.
- [24] Y. AlHumaidan, L. AlAjmi, M. Aljamea, and M. Mahmud, “Analysis of cloud computing security in perspective of Saudi Arabia,” in *2018 IEEE 20th International Conference on e-Health Networking, Applications and Services (Healthcom)*, IEEE, 2018, pp. 1–4. DOI:10.1109/HealthCom.2018.8531141.
- [25] D. Chen and H. Zhao, “Data security and privacy protection issues in cloud computing,” in *2012 international conference on computer science and electronics engineering*, IEEE, 2012, pp. 647–651. DOI:10.1109/ICCSEE.2012.193.
- [26] A. Khayer, M. S. Talukder, Y. Bao, and M. N. Hossain, “Cloud computing adoption and its impact on SMEs’ performance for cloud supported operations: A dual-stage analytical approach,” *Technol Soc*, vol. 60, p. 101225, 2020. DOI:10.1016/j.techsoc.2019.101225.
- [27] N. Al Mudawi, N. Beloff, and M. White, “Developing a framework of critical factors affecting the adoption of cloud computing in government systems (ACCE-GOV),” in *Intelligent Computing: Proceedings of the 2021 Computing Conference, Volume 1*, Springer, 2022, pp. 520–538. https://doi.org/10.1007/978-3-030-80119-9_32.
- [28] F. Alharbi, A. Atkins, and C. Stanier, “Understanding the determinants of Cloud Computing adoption in Saudi healthcare organisations,” *Complex & Intelligent Systems*, vol. 2, pp. 155–171, 2016. DOI:10.1007/s40747-016-0021-9.
- [29] H. Hassan, M. H. Mohd Nasir, and N. Khairudin, “Cloud computing adoption in organisations: review of empirical literature,” in *SHS Web of Conferences*, EDP Sciences., 2017. <https://doi.org/10.1051/shsconf/20173402001>.
- [30] H. Hassan, M. H. M. Nasir, N. Khairudin, and I. Adon, “Factors influencing cloud computing adoption in small medium enterprises,” *Journal of Information and Communication Technology*, vol. 16, no. 1, pp. 21–41, 2017.
- [31] A. Tarhini, A. Al-Badi, M. Almajali, and S. H. Alrabayah, “Factors Influencing Employees’ Intention to Use Cloud Computing,” *Journal of Management and Strategy*, vol. 8, no. 2, pp. 47–62, 2017.
- [32] A. Cardoso, F. Moreira, and D. F. Escudero, “Information technology infrastructure library and the migration to cloud computing,” *Univers Access Inf Soc*, vol. 17, no. 3, pp. 503–515, 2018.
- [33] M. Al-Ruithe, E. Benkhelifa, and K. Hameed, “Key issues for embracing the cloud computing to adopt a digital transformation: A study of Saudi public sector,” *Procedia Comput Sci*, vol. 130, pp. 1037–1043, 2018. <https://doi.org/10.1016/j.procs.2018.04.145>.
- [34] T. Alsafi and I.-S. Fan, “Investigation of Cloud Computing Barriers: A Case Study in Saudi Arabian SMEs,” *Journal of Information Systems Engineering and Management*, vol. 5, no. 4, p. em0129, 2020. <https://doi.org/10.29333/jisem/8534>.
- [35] C. Low, Y. Chen, and M. Wu, “Understanding the determinants of cloud computing adoption,” *Industrial management & data systems*, vol. 111, no. 7, pp. 1006–1023, 2011. DOI:10.1108/02635571111161262.
- [36] K. K. Hiran, “Impact of Driving Factors on Cloud Computing Adoption in the Higher Education,” in *IOP Conference Series: Materials Science and Engineering*, IOP Publishing, 2021, p. 012016. DOI:10.1088/1757-899X/1131/1/012016.
- [37] N. Alkhater, R. Walters, and G. Wills, “An empirical study of factors influencing cloud adoption among private sector organisations,” *Telematics and Informatics*, vol. 35, no. 1, pp. 38–54, Apr. 2018, doi: 10.1016/J.TELE.2017.09.017.
- [38] H. M. Sabi, F.-M. E. Uzoka, and S. V Mlay, “Staff perception towards cloud computing adoption at universities in a developing country,” *Educ Inf Technol (Dordr)*, vol. 23, pp. 1825–1848, 2018.
- [39] M. Jaradat, H. T. Ababneh, K. M. Faqih, and N. M. Nusairat, “Exploring cloud computing adoption in higher educational environment: an extension of the UTAUT model with trust,” *International Journal of Advanced Science and Technology*, vol. 29, no. 5, pp. 8282–8306, 2020.

A Quantitative Study Using the ACC-PH Framework: Factors Affecting Cloud Computing Adoption in Saudi Private Hospitals

Fayez Alshahrani
0009-0001-7389-0548
Department of Information Systems, Najran University, Najran, Saudi Arabia
Department of Informatics, University of Sussex, Brighton, United Kingdom
Email: fa461@sussex.ac.uk

Natalia Beloff
0000-0002-8872-7786
Department of Informatics, University of Sussex, Brighton, United Kingdom
Email: n.beloff@sussex.ac.uk

Martin White
0000-0001-8686-2274
Department of Informatics, University of Sussex, Brighton, United Kingdom
Email: m.white@sussex.ac.uk

Abstract—Private hospitals aim to provide essential healthcare services while focusing on profit and income growth. They are turning to innovative solutions to enhance medical services efficiency while reducing costs. Cloud computing has arisen as an ideal option, allowing private hospitals to access advanced digital health services without heavy infrastructure investments. Yet, in Saudi private hospitals, the adoption of Cloud computing is remarkably low. Therefore, in this study, we surveyed 650 managers and administrative staff from Saudi private hospitals, using our previously proposed ACC-PH framework to assess factors influencing Cloud computing adoption from technological, organisational, and environmental perspectives. The data were analysed using IBM-SPSS and AMOSv29. The results revealed the positive influence of 12 out of 13 examined factors. The findings are significant in guiding decision-makers in Saudi private hospitals to establish effective strategies for implementing Cloud computing. These strategies can enable easier adoption of Cloud computing in this essential industry.

Index Terms—Cloud computing adoption, Saudi Arabia, private hospitals.

I. INTRODUCTION

THE private healthcare industry in Saudi Arabia, especially private hospitals, is encountering significant challenges. Recent statistics from Saudi's Ministry of Health indicate that approximately 43% of the country's population receives healthcare services from private hospitals [1]. However, despite serving a large segment of the population, Saudi private hospitals grapple with the challenge of maintaining profitable returns while simultaneously providing high-quality medical services [2]. Therefore, Saudi private hospitals should consider creative solutions to balance delivering advanced healthcare and meeting their economic objectives. This pressing need for balance underscores the potential benefits of incorporating technological solutions, such as Electronic Health (E-Health) systems.

E-Health systems play a pivotal role in elevating the quality of healthcare services in private hospitals, but they also raise

economic challenges. E-Health systems can streamline patient data management, enhance diagnostic accuracy, and enable more efficient treatment processes [3], [4]. However, the implementation and maintenance of E-Health systems can be financially demanding. The initial investment in technological infrastructure, ongoing operational costs, and the need for specialised technicians to manage these systems contribute to increased expenses [5], [6]. As a result, this can negatively impact the profitability of private hospitals. Consequently, balancing the benefits of improved healthcare quality with the financial implications remains an enormous hurdle for private healthcare institutions. Therefore, cost-effective, innovative technologies such as Cloud computing can ease this hurdle.

Cloud computing offers an innovative solution that boosts private hospitals' performance efficiency at a lower cost. Cloud computing revolves around utilising electronic services from external providers, eliminating the necessity for costly technological infrastructure, maintenance expenses and in-house technical experts [7]. Hence, Cloud computing plays a crucial role in cost reduction and boosting the profitability of private hospitals. In addition to cost-effectiveness, Cloud computing offers superior E-health services compared to traditional E-health solutions. That is due to its greater scalability and enhanced ability to manage large volumes of health data [8]. Thus, Cloud computing can significantly increase the effectiveness of E-health services, leading to improved healthcare service quality. However, does Cloud computing also raise some associated concerns?

Yes, Cloud computing may present data security challenges, yet recommended strategies are available to address these issues. According to [9], patient data stored in the Cloud is considered less secure. The primary concern threatening patient data in the Cloud is unauthorised access, as the data are exposed to various platforms [10]. However, to address these

concerns, [11] suggests implementing robust security protocols, which can safely leverage the benefits of Cloud computing in E-Health systems.

Implementing these robust security protocols mitigates risks and paves the way for Cloud computing to revolutionise healthcare delivery. Therefore, with Cloud computing succeeding in addressing high-quality versus low-cost challenges, the ability to provide optimal healthcare services at reduced costs is now within the reach of Saudi private hospitals.

Nevertheless, despite these promising developments, the adoption of Cloud technology in Saudi private hospitals is still behind. The literature shows that among various sectors in Saudi Arabia, the healthcare sector, involving private hospitals, has the lowest usage of Cloud computing [12]. It is a phenomenon that highlights the necessity to investigate the affecting factors.

Prior research has shed light on factors that affect the adoption of Cloud computing across healthcare institutions in Saudi Arabia [13], [14], [15], [16]. Although previous studies reveal important effective factors, there is a gap in the literature that exclusively addresses the Saudi private healthcare sector, particularly private hospitals.

Therefore, focusing research on Saudi private hospitals' unique dual goals of quality healthcare service delivery and economic growth can yield distinctive insights not seen in the public sector. Such targeted research could produce more relevant findings that are vital for enabling decision-makers in these institutions to adopt Cloud computing.

Consequently, our paper aims to assist Saudi private hospital decision-makers in formulating suitable strategies for implementing Cloud computing by determining adoption influencing factors. Initially, the paper presents a review of relevant literature and outlines the foundational framework of this study. Following this, the paper delves into the data collection and analysis methodology. The research then provides details on the study's findings, leading to a discussion section that evaluates these results within the broader research field. Finally, in the conclusion, the paper lays the groundwork for our related future research work.

II. LITERATURE REVIEW

In Saudi Arabia, studies have investigated the impact of factors from various contexts on adopting Cloud computing in the healthcare sector. A study by [16] targeted Saudi public and private healthcare institutions to investigate technological, organisational, environmental, human, and business factors. Another study targeted the same population but only investigated factors from a technological context, which was carried out by [14]. The third study in the research field focused on Saudi university hospitals in Riyadh City, investigating the influence of the contexts of technology, organisation, environment, and decision-making on the adoption of Cloud computing [13]. The last and most recent study targeted Saudi hospitals without specifying public or private, focusing on influences on the adoption of Cloud computing

from human, organisational, environmental, and technological contexts [15]. A clear gap can be identified after reviewing the related studies in the literature.

The scarcity of these studies indicates a notable lack of research, specifically research that exclusively targets the private healthcare sector in Saudi Arabia. Traditionally, research may have concentrated more on public institutions due to their accessibility, public funding transparency, and the broader impact on national health policy. That often leaves private sectors less inspected. The lack of specific focus on private hospitals is significant, as these entities frequently implement technology differently than public institutions, potentially leading to varied adoption rates and challenges. That is a substantial gap in the literature which needs to be filled. It underscores the urgency for targeted research; fortunately, the existing studies provide a robust foundation that can guide in-depth exploration of how healthcare institutions in Saudi Arabia can effectively adopt Cloud computing technologies.

Based on these foundations, four key technological factors have emerged as particularly influential in the adoption of Cloud systems within Saudi's healthcare sector: Compatibility, as defined by studies [13], [14], [15], assesses how well the prospective technology aligns with an organisation's existing values and operations. Although not a focus of the study [16], this factor showed a substantial impact in other studies. Security, which relates to the effectiveness of Cloud providers in protecting data, was notably emphasised in studies [14] and [15]. Relative Advantage evaluates the additional benefits a hospital gains through Cloud computing and was highlighted as significant in studies [13] and [16]. Reliability, which involves the system's ability to perform anticipated tasks effectively and securely, was exclusively studied in [14] and found to have a considerable impact.

In addition, factors including Top Management Support, Prior Experience, Organizational Readiness, Attitude towards Change, and Cost Analysis were found in previous studies to be the most impactful from an organisational perspective. Top Management Support was explored in studies [13], [15], [16] with findings indicating that leadership significantly drives Cloud adoption by allocating necessary resources. Also, research [13] showed that the prior technological experiences of top managers in Saudi healthcare organisations (Prior Experience) significantly influenced their decisions to migrate to Cloud-based systems. Moreover, the study [15] noted the critical role of Organizational Readiness, referring to the role of human, technological, and financial resources in facilitating Cloud adoption.

Furthermore, Attitude towards Change and Cost Analysis are key organisational factors confirmed by studies [14], [16]. While the study [14] examined Attitude towards Change under technological factors, we argue that since this factor assesses the impact of employees' beliefs about Cloud computing on adoption, it should be regarded as part of the organisational context, considering employees as human resources. Similarly, the Cost Analysis factor demonstrated its signifi-

cant impact in the study [16] within a business context. It involves a comprehensive cost versus-benefits analysis of Cloud adoption that influences decision-making. We emphasise that the Cost Analysis factor, as performed by relevant employees within healthcare organisations, should also be classified under the organisational context, underscoring its role in shaping strategic decisions.

The literature review offered insights into these five factors that fall within the organisational context and affect the adoption of Cloud computing in healthcare facilities in Saudi Arabia.

The literature also identified other influencing factors from the environment surrounding Saudi healthcare organisations. One was Competitive Pressure, where decision-makers are urged to match competitors already adopting Cloud computing, as noted in the study [15]. Although the study [13] also examined this factor, it did not find it significantly impactful. Another critical environmental influencer is Rules and Regulations, which was noted considerably in the study [16] as government policies play a decisive role in healthcare centres' decisions to adopt Cloud technologies, unlike in studies [13], [15] where it wasn't as pronounced.

These environmental factors and previously identified technological and organisational ones provide a comprehensive foundation for a new model to evaluate the determinants of Cloud computing adoption in Saudi private hospitals. Building on this, we have developed the ACC-PH framework [17], incorporating these eleven core factors along with two additional critical elements pertinent to technology adoption in developing countries: the availability of Cloud providers and Internet connectivity, which are scarce [18] and limited [19] in Saudi Arabia, respectively. This model consolidates thirteen factors hypothesised to positively influence Cloud adoption in private healthcare settings, as illustrated in Fig. 1.

The ACC-PH model gains its robustness from its foundation on two leading theories of technology adoption: The Model of Diffusion of Innovation (DOI) and The Theory of Technology, Organisation, and Environment (TOE). The ACC-PH theory incorporates two factors from the DOI theory: Relative Advantage and Compatibility. The ACC-PH framework also categorises its factors into the three contexts outlined by TOE: technological, organisational, and environmental. Therefore, the ACC-PH framework is predicted to have a significant theoretical contribution to the research field and practical application in assisting Saudi private hospitals in gaining the benefits of adopting Cloud computing.

To ensure that the theoretical insights of the ACC-PH are effectively translated into practical research actions, the next section will detail the specific approaches and procedures used in the study application. The following methodology section will provide a comprehensive overview of the data collection and analysis techniques for applying the ACC-PH model to our study.

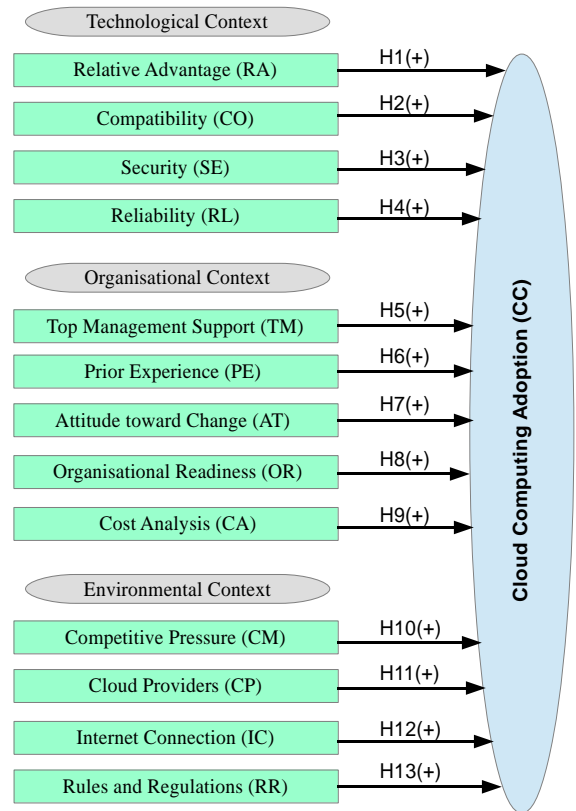


Fig. 1 The Comprehensive Framework for Adopting Cloud Computing in Private Hospitals (ACC-PH) Adapted From [17]

III. METHODOLOGY

The research employs a scientific method to thoroughly assess the ACC-PH Framework's suggested factors influencing Cloud computing adoption in Saudi private hospitals. The employed methodology allowed for an efficient exploration of how these factors impact the decision-making processes in healthcare settings.

The approach started with determining the target population. The study targeted managers and administrative staff to ensure a comprehensive evaluation of Cloud computing adoption in Saudi private hospitals. They were specifically chosen as the target population for this study because they are typically the decision-makers regarding implementing new technologies within hospitals. This strategic selection ensures that the survey captures insights from those directly involved in the decision-making process concerning technological adoption. Having identified the key stakeholders for our survey, we next determined the necessary sample size to ensure reliable results.

A standard sample size formula, developed by statisticians, was applied to calculate the sample size, considering the large yet unspecified population size. The available statistics from the Saudi Ministry of Health primarily focused on healthcare professionals like doctors, nurses, and other medical practitioners, without specific details on administrative staff. However, these statistics also revealed that until 2021, there are

approximately 115 accredited private hospitals in Saudi Arabia, indicating a substantial population size of administrative employee numbers [1]. Thus, the formula to compute the standard sample size used a confidence level of 95% ($Z = 1.96$), an error margin of 5% ($E = 0.05$), and an assumed population proportion, $p = 0.5$, to ensure the sample size maximisation. The result disclosed the minimum representative sample of administrative staff working in Saudi private hospitals is approximately 384 individuals. As the sample size was firmly established, we carefully crafted the questionnaire to evaluate the various factors influencing Cloud computing adoption.

Comprehensive and transparent questionnaire questions were designed to assess each proposed factor. Four questions (items) were assigned to each factor (latent construct) to prevent the expected exclusion of certain items from affecting the evaluation of the factor in the analysis stage. Various questions were also considered to assess each factor from multiple perspectives. In addition, the survey questions were structured using a 5-point Likert scale, with one representing "Strongly disagree" and five denoting "Strongly agree". After finalising the questionnaire, the focus was on selecting advanced tools for efficient data collection and analysis.

The research methodology included identifying tools for data collection and analysis. Surveys were designed and administered using Qualtrics Core XM software for online data collection. Qualtrics was selected due to its capability to create tailored evaluation instruments and enhance the efficiency of data collection [20]. The surveys were distributed via email, and the collected data were saved in password-protected files on Microsoft OneDrive. In the data analysis phase, we used IBM-SPSS and AMOS version 29. IBM-SPSS and AMOS were explicitly chosen for their advanced capabilities in structural equation modelling, which is essential for effectively analysing the complex interrelations inherent in the comprehensive ACC-PH model. These tools enable precise testing and validation of the model's constructs, ensuring a robust analysis of the data collected [21]. With the data collection and analysis tools set, the next critical step was ensuring all research aspects adhered to the highest ethical standards.

The research was granted ethical clearance by the Sciences & Technology Cross-Schools Research Ethics Committee (C-REC) at the University of Sussex. This approval is documented by the ethical review application number

[ER/FA461/2]. The ethical approval came with a guarantee of anonymity and confidentiality for all participants. Participants were assured that their information would be used exclusively for academic purposes and remain confidential. Additionally, the study was designed to ensure that no third parties are involved and that participants' data will not be shared with any external entities. To further protect participant identity, each individual was not asked for personal information but instead assigned a unique identifier, distinguishing their participation while maintaining the anonymity of their contributions.

Following the stringent adherence to ethical standards, the study was well-positioned to proceed into the data analysis phase. The results, derived from the gathered data, are detailed in the following section.

IV. RESULTS

This section presents the study's comprehensive findings. We present the outcomes of the investigation, starting with an exposition of demographic statistical data. It is followed by an evaluation of the construct's reliability and validity. Subsequently, we illustrate the findings derived from the empirical examination of the proposed hypotheses within the ACC-PH framework.

A. Demographic Statistics

Demographic statistics are essential for validating research findings. Describing the characteristics of study participants plays a crucial role in enabling generalisability and ensuring representativeness, thereby improving the overall quality of the research [22]. Demographic statistics provide data like age, gender, education level, and other demographics. Table I. details the demographic results of our study. These details offer a comprehensive breakdown of participant demographics to elucidate the representativeness and diversity observed within the study.

The demographic composition of the 650 participants from Saudi private hospitals surveyed between 15th June 2023 and 15th September 2023 offers insights into the diverse backgrounds represented in the study, including job roles, gender, age, and educational levels. Of these participants, 396 (60.9%) were IT employees, 227 (34.9%) held other administrative roles, and 27 (4.2%) were directors and top managers. Among the participants, 276 (42.5%) were male and 374

TABLE I.
DEMOGRAPHIC STATISTICS FOR THE PARTICIPANTS

Statistic Type	Category	Frequency & Percentage	Statistic Type	Category	Frequency & Percentage
Job Title	IT employees	396 (60.9%)	Gender	Male	276 (42.5%)
	Directors	27 (4.2%)		Female	374 (57.5%)
	Other administrative	227 (34.9%)		Secondary or less	10 (1.5%)
Age	18-35	468 (72%)	Level of Education	Diploma	66 (10.2%)
	36-45	118 (18.2%)		Bachelor	449 (69.1%)
	46-60	56 (8.6%)		Master	118 (18.1%)
	Over 60	8 (1.2%)		PhD	7 (1.1%)

(57.5%) were female. The age distribution showed a significant majority of younger participants, with 468 (72%) aged between 18 and 35, while 182 (28%) were 36 and above. Educationally, 449 (69.1%) of the participants held a bachelor's degree, followed by 118 (18.1%) with master's degrees, 66 (10.2%) with diplomas, 10 (1.5%) with a secondary school diploma or less, and 7 (1.1%) with Doctoral degrees.

After delineating the demographic profile of the participants, attention was subsequently directed toward verifying the statistical prerequisites necessary for employing Structural Equation Modelling (SEM) in our analysis.

B. Normality Distribution

As we utilised Structural Equation Modelling (SEM) to analyse the relationship between observed variables (items) and latent constructs (factors), ensuring the normal distribution of the observed data was crucial. SEM is a parametric statistical technique that yields efficient and reliable estimates when the data are normally distributed [23]. We checked univariate and multivariate outliers, skewness, kurtosis, and standard deviations to verify normality distribution.

Univariate outliers, extreme data points within individual variables, were detected using z-scores; values beyond ± 3.29 are considered outliers. Multivariate outliers were identified across multiple variables using Mahalanobis Distance, with a p-value lower than 0.05 indicative of outliers [24]. The results indicated no univariate outliers, as the maximum z-score was 1.77, and the minimum was -3.11. Similarly, no multivariate outliers were found, as the maximum observed Mahalanobis D-squared value was 23.798, corresponding to a p-value of 0.051.

Additionally, skewness and kurtosis values fell within the acceptable range of ± 2.58 , suggested by [25], with the lowest skewness at -0.755, the highest at 0.010, and the maximum and minimum kurtosis values at 0.531 and -0.796, respectively. The standard deviations for all observed variables were close to zero, with a maximum deviation of 1.11 and a minimum of 0.91.

These findings confirmed the normal distribution of the observed data, which was necessary for the SEM analysis of the proposed complex ACC-PH model.

Given the confirmation of data normality, the subsequent analysis moves to examine the model's structural integrity. Specifically, the following section discusses Confirmatory Factor Analysis (CFA), an essential Structural Equation Modelling (SEM) step.

C. The Confirmatory Factor Analysis (CFA)

Confirmatory Factor Analysis (CFA) is a crucial initial step in testing hypotheses using the Structural Equation Modelling (SEM) approach. CFA validates proposed models by analysing the interactions between latent constructs and observed variables and provides metrics like standardised factor loadings and latent variable correlations for evaluating the framework's applicability [26]. Therefore, the CFA for our proposed ACC-PH model was designed using AMOS version 29. Thus, with the CFA rigorously set up through AMOS version

29, the following section will detail the outcomes of the measures computed in the CFA, evaluating how well the latent constructs and observed variables align within the ACC-PH model.

D. Unidimensionality Assessment (Construct Validity)

Unidimensionality assessment in CFA involves evaluating whether a set of observed variables accurately and appropriately represents a single underlying latent construct, which is essential for maintaining the coherence of measurements related to latent constructs [27]. This verification process is crucial in increasing confidence in the results as the first step in ensuring construct validity [28].

Various methods exist to assess unidimensionality, but standardised factor loading (λ) analysis remains fundamental in confirming that observed variables are sufficient representatives of their respective construct. There is some debate over the standardised factor loading threshold for assessing unidimensionality, with some scholars advocating for a 0.40 threshold (for exploratory studies) [29] while others suggest a stricter 0.60 criterion [28]. This research adopted a more rigorous 0.60 threshold as the standard criterion to accept or reject observed variables.

In implementing these criteria, 14 latent constructs were measured using 56 observed variables in this study, with 55 displaying a standardised loading factor (λ) above 0.60. However, one variable representing the Security (SE) latent construct did not meet the standard criterion with a standardised loading factor (λ) of 0.59 and was therefore omitted from the CFA. This ensured that all remaining variables coherently represent their respective constructs, thereby preserving construct validity. This meticulous approach to variable selection helps preserve construct validity, setting the stage for further validation through additional measures such as Model Fit Indices in CFA.

E. Model Fit Metrics (Construct Validity)

Confirmatory Factor Analysis (CFA) is used to further evaluate the construct validity of a measurement model through various model fit metrics/indices. Critical tools such as the Standardised Root Mean Square Residual (SRMR), Chi-Square divided by degrees of freedom (CMIN/DF), Incremental Fit Index (IFI), Tucker-Lewis Index (TLI), Comparative Fit Index (CFI), Root Mean Square Error of Approximation (RMSEA) along with its P-Close value, and Goodness of Fit Index (GFI) provide quantitative assessments of how well the model corresponds to the observed data [23], [30]. Each index offers a unique perspective on model fit, addressing different aspects such as discrepancy, incremental fit, and parsimony.

However, interpreting these indices requires caution due to their inherent limitations and potential to yield contradictory outcomes. For instance, indices like CFI and RMSEA are known for their reduced sensitivity to sample size variations, whereas CMIN/DF may be affected by larger sample sizes [23].

TABLE II.
MODEL FIT METRICS

Metric	SRMR	CMIN/DF	IFI	TLI	CFI	RMSEA	P-Close	GFI
Estimate	0.0336	1.372	0.962	0.958	0.962	0.024	1.000	0.901
Threshold	< 0.09 [35]	< 3.0 [25]	> 0.90 [36]	> 0.90 [37]	> 0.90 [37]	< 0.05 [30]	> 0.05 [38]	> 0.90 [39]
Evaluation	Good	Good	Good	Good	Good	Good	Good	Good

Therefore, acknowledging the complexity of accurately capturing a model's fit with a single index is crucial [23]. Thus, a comprehensive evaluation of construct validity necessitates the integration of multiple indices. This study assessed multiple model fit metrics, all remaining within the recommended thresholds, as detailed in Table II.

With the model's fit to the observed data confirmed by multiple indices, the next critical step involves scrutinising Construct Reliability to verify the stability and consistency of our measurement constructs.

F. Construct Reliability

Before further evaluations of construct validity, it is essential to ensure the measurement tool produces stable and consistent results by examining construct reliability. To achieve elevated levels of internal consistency, the variables on a scale must reliably measure the same construct. Specifically, for a construct to be deemed reliable, responses must remain consistent over time [25].

Construct reliability is typically evaluated using two main methods: Cronbach's Alpha (α) and Composite Reliability (CR) [31]. As shown in Table III., the results from these assessments confirmed that both Cronbach's Alpha values and Composite Reliability values exceeded the 0.70 threshold, aligning with [29]'s standards. According to these standards, scores above 0.70 in both Cronbach's Alpha and Composite Reliability indicate adequate construct reliability.

Following the confirmation of construct reliability through Cronbach's Alpha and Composite Reliability, the analysis progresses to assessing Convergent Validity, a critical aspect

of Construct Validity, to further validate the coherence and relevance of the constructs within the model.

G. Convergent Validity (Construct Validity)

Convergent validity is a crucial measurement that assesses whether different items or observed variables intended to measure the same construct yield similar outcomes. It aims to confirm that each item accurately measures its intended construct by demonstrating strong correlations with other items that target the same construct [25].

Methods such as the Average Variance Extracted (AVE) and Composite Reliability (CR) are employed to evaluate convergent validity. According to [29], an AVE value of 0.50 or higher indicates good convergent validity. Furthermore, [32] suggests that convergent validity is robust if the CR value exceeds the AVE value, implying that a significant portion of the variance in the items stems from the construct itself rather than from random measurement error.

In this study, as detailed in Table III, the constructs met these stringent criteria, with AVE values exceeding 0.50 and CR values surpassing AVE, thereby affirming the convergent validity of the constructs analysed.

As convergent validity ensures consistency among related items, it is equally crucial to establish the distinctiveness of unrelated items, a measurement to be explored next through Discriminant Validity.

H. Discriminant Validity (Construct Validity)

Discriminant validity is vital for establishing the distinctiveness of constructs within Confirmatory Factor Analysis

TABLE III.
CONSTRUCT RELIABILITY, CONVERGENT VALIDITY, AND DISCRIMINANT VALIDITY

Constructs	α	CR	AVE	MSV	Evaluation
Relative Advantage (RA)	0.869	0.872	0.526	0.081	Construct Reliability $\alpha > 0.70$ [29]
Compatibility (CO)	0.798	0.800	0.633	0.023	verified
Security (SE)	0.797	0.799	0.504	0.037	Construct Reliability CR > 0.70 [29]
Reliability (RL)	0.809	0.810	0.571	0.055	verified
Top Management Support (TM)	0.876	0.877	0.517	0.012	
Prior Experience (PE)	0.806	0.806	0.641	0.023	Convergent Validity AVE ≥ 0.5 [29]
Attitude toward Change (AT)	0.802	0.803	0.511	0.086	verified
Organisational Readiness (OR)	0.800	0.800	0.506	0.109	Convergent Validity CR > AVE [32]
Cost Analysis (CA)	0.806	0.806	0.502	0.080	verified
Competitive Pressure (CM)	0.813	0.814	0.511	0.102	
Cloud Providers (CP)	0.811	0.812	0.524	0.084	Discriminant Validity AVE > MSV [32]
Internet Connection (IC)	0.802	0.802	0.520	0.102	verified
Rules and Regulations (RR)	0.815	0.815	0.504	0.081	

(CFA). This form of validity requires that items from unrelated constructs exhibit either no or very low correlations, ensuring that each construct is uniquely captured [25]. There are two primary methods employed to assess discriminant validity.

The first method compares the Average Variance Extracted (AVE) against the Maximum Shared Variance (MSV). To establish discriminant validity, the AVE values for each construct must exceed the corresponding MSV values [32]. Table III. shows that for all latent constructs, the AVE values exceeded the MSV values, thus confirming discriminant validity according to this criterion.

The second method utilises the Fornell-Larcker Criterion, which requires the AVE's square root for each factor to be compared with the correlations or covariances between that factor and others within the model. According to this criterion, discriminant validity is confirmed when the AVE's square root for each factor exceeds any correlation with other factors [33]. The Fornell-Larcker Criterion is typically represented in a correlation matrix, where square roots of AVEs are displayed along the diagonal, and the off-diagonal values denote correlations between constructs. The outcomes of this test, as shown in Table IV, also met the recommended threshold criteria, further validating the discriminant validity of the constructs.

With all measures of Confirmatory Factor Analysis (CFA) affirming the reliability and validity of the constructs, the study now transitions to the next critical phase: building the Structural Equation Modelling (SEM) to thoroughly test the hypotheses.

I. Structural Equation Modelling (SEM)

After completing the Confirmatory Factor Analysis (CFA) for the ACC-PH model, we progressed to developing the Structural Equation Model (SEM). SEM is a renowned multivariate quantitative approach widely utilised for hypothesis

testing and offers significant benefits due to its comprehensive nature. This method facilitates the parallel analysis of multiple relationships and effectively handles latent constructs [23].

Consequently, the SEM was constructed to include 13 independent constructs: Relative Advantage (RA), Compatibility (CO), Security (SE), Reliability (RL), Top Management Support (TM), Prior Experience (PE), Attitude towards Change (AT), Organisational Readiness (OR), Cost Analysis (CA), Competitive Pressure (CM), Cloud Providers (CP), Internet Connection (IC), and Rules and Regulations (RR). Cloud Computing Adoption (CC) was designated as the model's dependent construct, central to integrating these various factors.

After establishing the Structural Equation Model (SEM), the next phase of our research focuses on the findings derived from applying SEM to test the hypotheses.

J. Results of Testing Hypotheses within the ACC-PH Framework

In the Structural Equation Modelling (SEM) applied for this study, four key metrics were computed to test the hypotheses of the ACC-PH model: Standardised Regression Weight (β), Standard Error (SE), Critical Ratio (CR), and P-value. These metrics determine the significance of the relationships between the independent variable and the dependent variables.

According to [23], Standardised Regression Weight (β) measures the intensity and direction of a variable's association; larger absolute numbers indicate a stronger connection. Positive values suggest a positive relationship, while negative values indicate a negative relationship. The Standard Error (SE) estimates the variability of the regression weight, with lower values indicating more precision and reliability. The Critical Ratio (CR) is calculated by dividing the regression weight by the standard error, with values beyond ± 1.96 typi-

TABLE IV.
DISCRIMINANT VALIDITY (FORNELL-LARCKER CRITERION)

	RR	RA	CO	SE	RL	TM	PE	AT	OR	CA	CM	CP	IC	CC
RR	0.725													
RA	0.067	0.796												
CO	0.081	0.001	0.710											
SE	0.156	0.153	0.055	0.756										
RL	0.044	-0.030	0.059	-0.001	0.719									
TM	0.107	-0.007	0.015	-0.036	0.002	0.801								
PE	0.213	0.060	0.169	0.187	-0.110	0.014	0.715							
AT	0.044	-0.031	0.112	0.030	-0.068	0.008	0.117	0.711						
OR	0.137	0.022	0.084	0.085	0.035	0.009	0.249	0.245	0.708					
CA	0.165	0.014	0.004	0.091	0.045	-0.042	0.068	0.108	0.164	0.715				
CM	0.281	0.005	-0.057	0.076	0.032	0.113	0.290	-0.137	0.155	0.074	0.724			
CP	0.114	0.114	0.156	0.093	-0.051	-0.050	0.058	-0.018	-0.034	0.320	0.034	0.721		
IC	0.285	0.113	0.007	0.071	0.088	-0.018	0.140	0.020	0.102	0.177	-0.071	0.275	0.710	
CC	0.281	0.152	0.193	0.234	0.091	0.150	0.293	0.330	0.282	0.248	0.185	0.220	0.282	0.716

cally denoting statistical significance at the p-value significant level. The P-value assesses the probability of the observed effect occurring by chance under the null hypothesis, where a value below 0.05 typically denotes statistical significance, suggesting a non-random effect.

Based on these methods, the hypotheses listed in Table V. were tested, with results shown in Table VI. This study assessed four technological factors influencing Cloud computing adoption in Saudi Arabian private hospitals. Security was identified as the most critical factor, followed by Compatibility, Reliability, and Relative Advantage. Additionally, five organisational factors positively influenced Cloud computing adoption, with the employees' Attitude towards Change being the most impactful, followed by Top Management Support. Other significant factors included Prior Experiences, Organisational Readiness, and Cost Analysis.

Furthermore, environmental factors such as Internet Connection, Competitive Pressure, and Cloud Providers were analysed and determined to influence adoption substantially. However, Rules and Regulations have not significantly influenced the adoption of Cloud computing.

Upon detailing the study's findings, engaging in a discussion contextualised within the literature review and a pragmatic framework becomes imperative. The following discussion will ensure that our research's theoretical and practical contributions are achieved.

V. DISCUSSION

Our study investigates several factors influencing the adoption of Cloud computing in private hospitals across Saudi Arabia, focusing on technological, organisational, and environmental aspects. The findings not only corroborate existing theoretical frameworks but also unveil unique insights pertinent to the specific challenges within the Saudi private healthcare sector. By analysing each factor's distinct contributions and their interrelationships, this research provides a nuanced understanding of the dynamics involved in adopting

Cloud computing technologies within Saudi Arabia's private healthcare settings.

A. Technological Context

The study investigated the impact of various technological factors on the adoption of Cloud computing in Saudi private hospitals. The findings revealed a positive influence of these factors on Cloud computing adoption, supported by empirical evidence.

Security (SE) was identified as the most influential technological factor. This was indicated by its notable Standardised Regression Weight (β) of 0.121 and the statistically significant P-value of 0.003. The findings align with existing scientific research, such as studies [14], [15], which have similarly highlighted the positive role of Security in adopting Cloud computing within the broader Saudi healthcare context. This agreement underscores the importance of robust security measures, which are crucial for successfully integrating Cloud computing in healthcare environments, especially in private hospitals in Saudi Arabia. The emphasis on security reflects a broader trend in healthcare technology, where protecting sensitive patient data and maintaining system integrity is paramount. It is also essential to consider advanced cyber threats, such as ransomware and advanced persistent threats (APTs), which pose significant risks to Cloud technologies in the health sector [34].

Following Security, **Compatibility (CO)** appeared as the second most significant factor, with a Standardised Regression Weight (β) of 0.113 and a P-value of 0.018. This finding is corroborated by studies [13], [14], [15], which have previously drawn similar conclusions. The significance of Compatibility lies in its facilitation of seamless integration and effective use of Cloud computing technologies within existing healthcare systems. Ensuring that new technologies align well with healthcare providers' current practices, workflows, and needs is particularly vital. High levels of compatibility enhance the likelihood of successful adoption and implementation of Cloud computing in Saudi private hospitals.

TABLE V.
LIST OF HYPOTHESES WITHIN THE ACC-PH FRAMEWORK AND THEIR RESULTS ADAPTED FROM [17]

H#	Hypothesis Statement	Result
H1	"Recognising the relative advantage of Cloud technology enhances the likelihood of adopting Cloud computing in Saudi private hospitals."	Supported
H2	"Higher compatibility enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H3	"Higher security levels enhance the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H4	"Higher reliability enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H5	"Top management support enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H6	"Top managers' sufficient prior technical experience enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H7	"Positive employees' attitudes towards change enhance the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H8	"Organisational readiness enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H9	"Efficient cost analysis enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H10	"Competitive pressure enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H11	"Cloud providers' availability within the same country enhances the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H12	"Internet connection availability and high functionality enhance the likelihood of adopting cloud computing in Saudi private hospitals."	Supported
H13	"Flexible rules and regulations enhance the likelihood of adopting cloud computing in Saudi private hospitals."	Rejected

TABLE VI.
RESULTS OF ANALYSIS OF HYPOTHESIS PATH FOR THE ACC-PH FRAMEWORK

Context	Structural Relation	Regression Weight (β)	Standard Error (S.E.)	Critical Ratio (C.R.)	P-value
Technological	SE \rightarrow CC	0.121	0.041	2.938	0.003**
	CO \rightarrow CC	0.113	0.048	2.371	0.018*
	RL \rightarrow CC	0.096	0.044	2.188	0.029*
	RA \rightarrow CC	0.071	0.029	2.432	0.015*
Organisational	AT \rightarrow CC	0.310	0.052	5.905	< 0.001***
	TM \rightarrow CC	0.118	0.034	3.445	< 0.001***
	PE \rightarrow CC	0.112	0.048	2.323	0.020*
	OR \rightarrow CC	0.106	0.052	2.050	0.040*
	CA \rightarrow CC	0.093	0.045	2.065	0.039*
Environmental	IC \rightarrow CC	0.166	0.050	3.347	< 0.001***
	CM \rightarrow CC	0.124	0.050	2.500	0.012*
	CP \rightarrow CC	0.116	0.053	2.189	0.029*
	RR \rightarrow CC	0.069	0.050	1.382	0.167

The study also confirmed the core role of the **Reliability (RL)** factor, with a Standardised Regression Weight (β) of 0.096 and a P-value of 0.029. In the literature, the study [14] examined the impact of the RL factor in the Saudi healthcare sector and gave equivalent results. The emphasis on Reliability is significant in healthcare, where consistent, dependable, and uninterrupted E-Health services are non-negotiable. The healthcare sector's unique demands, such as the need for constant access to patient records and operational continuity, make Reliability a significant factor influencing the decision to adopt Cloud computing technologies in private hospitals in Saudi Arabia.

Lastly, the study identified the **Relative Advantage (RA)** as another significant factor, indicated by a Standardised Regression Weight (β) of 0.071 and a P-value of 0.015. The findings support the outcomes of most of the previous research in the Saudi healthcare sector [13], [16]. These results highlight the importance of recognising the benefits and added value of Cloud computing to private healthcare institutions. The focus on Relative Advantage is particularly relevant in strategic decision-making. Understanding the tangible and intangible benefits of Cloud computing can drive adoption. Our findings suggest that private healthcare facilities are more likely to adopt Cloud computing solutions when they perceive clear advantages over existing systems, such as improved efficiency, cost savings, scalability, and enhanced data management capabilities.

As the technological aspects are thoroughly evaluated, the focus now shifts to the organisational factors. These factors are key to understanding private hospitals' internal readiness to adopt Cloud technologies, highlighting the human and strategic drivers of technology integration.

B. Organisational Context

This research measured five organisational factors potentially impacting Cloud computing adoption in private-sector hospitals in Saudi Arabia, confirming their positive influence.

Attitude towards Change (AT) emerged as the top factor, with a substantial Standardised Regression Weight ($\beta = 0.310$) and a highly significant P-value (less than 0.001). These results align with findings from studies [14], [16] although they diverge from [15]. The divergence likely stems from the multifaceted nature of technology adoption across different settings, influenced by cultural, structural, regulatory, and financial elements. In private hospitals, there is often a stronger emphasis on efficiency, innovation, and customer satisfaction due to the competitive nature of the healthcare market. Cultural factors such as the organisational culture towards innovation, structural factors like the flexibility of management structures, regulatory elements including compliance with local and international standards, and financial constraints or incentives play significant roles in shaping technology adoption. These diverse influences can lead to significant differences in how technology is implemented and utilised, explaining the divergent findings. Our study highlights the critical role of a workforce receptive to change, suggesting that adaptability and positive attitudes towards new technologies are significant drivers for successfully integrating Cloud computing in the Saudi private healthcare sector. These findings emphasise the need to cultivate a change-embracing organisational culture, which is vital in the evolving healthcare technology landscape.

Top Management Support (TM) was identified as the second most influential factor, with a Standardised Regression Weight ($\beta = 0.118$) and a P-value of less than 0.001, confirming its positive impact. These findings are consistent with previous studies [13], [15], indicating the importance of leadership and strategic guidance in adopting Cloud computing technologies in healthcare. Top management's support, resource allocation, and strategic vision are indispensable for implementing and efficiently utilising Cloud computing in Saudi Arabia's private healthcare facilities.

Prior Experiences (PE) of leaders and decision-makers in private hospitals in Saudi Arabia emerged as the third most influential factor, with a Standardised Regression Weight (β

= 0.112) and a P-value of 0.020. This finding aligns with the study [13] but contrasts with other literature, possibly due to the different sample populations studied. Our study's exclusive focus on Saudi private hospitals may explain these differences. Private hospitals often operate under various constraints and motivations compared to public hospitals. In private healthcare settings, decision-makers may value prior experience more because it directly impacts their ability to innovate and stay competitive in a market-driven environment. In contrast, public sector hospitals, which are often more bureaucratic and have different funding structures, may not prioritise prior experience to the same extent. This distinction highlights how hands-on experience and familiarity with similar technologies can positively affect the readiness and capability of decision-makers in private hospitals to integrate Cloud computing into their operational frameworks.

The study also highlighted **Organisational Readiness (OR)** as a significant factor in transitioning to Cloud computing, with a Standardised Regression Weight ($\beta = 0.106$) and a P-value of 0.040. These findings support the study [15] but contrast with [13]. The disparity in findings could be due to differences in the scope and focus of the studies, such as variations in sample size, research methodology, or regional healthcare policies. Our analysis suggests that an organisation's readiness, including its human, financial, and technological resources, is vital for successfully adopting Cloud computing in Saudi Arabia's private healthcare sector. Ensuring institutions are adequately prepared at multiple levels—from infrastructure to staff training and financial planning—is crucial for embracing and benefiting from Cloud computing solutions.

Finally, **Cost Analysis (CA)** was identified as a positive and direct determinant among organisational factors, with a Standardised Regression Weight ($\beta = 0.093$) and a significant P-value (0.039). This finding is consistent with existing literature, such as the study [16]. The results confirm that practical cost analysis influences Cloud computing adoption in Saudi private hospitals. Healthcare administrators and decision-makers need to understand the financial implications and potential advantages of Cloud computing enabling more informed and strategic choices.

Having explored the organisational influences on Cloud adoption, the discussion translates to environmental factors. These external elements play a critical role in shaping the adoption strategies of private hospitals in Saudi Arabia, highlighting the broader market and regulatory conditions that impact technological advancements.

C. Environmental Context

The analysis of quantitative data collected in this study revealed a significant and positive impact of environmental factors such as Internet Connection (IC), Competitive Pressure (CM), and Cloud Providers (CP) on the adoption of Cloud computing in Saudi private hospitals. Conversely, the Rules and Regulations (RR) factor did not significantly impact Cloud computing adoption. These findings contribute to understanding the environmental factors relevant when private

healthcare institutions in Saudi Arabia transition to Cloud systems.

Internet Connection (IC) emerged as a primary environmental factor that private hospital managers in Saudi Arabia need to consider during their transition to Cloud computing systems. Structural Equation Modelling (SEM) analysis showed a significant positive association between IC and Cloud Computing Adoption (CC), indicated by a Standardised Regression Weight (β) of 0.166 and a P-value of less than 0.001. These results underscore the criticality of IC in fostering Cloud computing adoption in Saudi private hospitals, emphasising the necessity of robust and reliable Internet connectivity. The complete dependence of Cloud services on the Internet, especially considering its restricted availability in a developing country like Saudi Arabia, highlights the significance of considering this factor in the adoption decision-making process. Moreover, exploring the IC factor provides novel insights into the literature, particularly in the Saudi healthcare sector, where such a relationship has not been previously investigated.

The influence of **Competitor Pressure (CM)** on Cloud computing adoption in Saudi private hospitals was also noteworthy, ranking second among environmental factors. SEM analysis confirmed a positive relationship between CM and CC, with a Standardised Regression Weight (β) of 0.124 and a P-value of 0.012. This finding shifts the balance in the literature regarding CM, supporting its positive impact on Cloud computing adoption in the Saudi healthcare sector and reconciling previous mixed results [13], [15]. The finding signifies that the impetus to remain competitive and technologically advanced is a substantial motivator for adopting Cloud computing among private hospitals in Saudi Arabia. This underscores the importance of these healthcare institutions adopting advanced technologies like Cloud computing to sustain competitiveness in the swiftly changing healthcare sector.

Furthermore, the study determined the **Cloud Providers (CP)** as another positive influential environmental determinant, ranking third. The SEM results validated this factor with a Standardised Regression Weight (β) of 0.116 and a P-value of 0.029. This finding reinforces the need for geographical proximity and the accessibility of Cloud service providers, suggesting that local availability significantly facilitates Cloud computing integration in Saudi Arabia's private healthcare sector. Additionally, investigating the CP factor enriches the existing literature by offering new perspectives in the context of the Saudi healthcare sector, where this factor has not been extensively explored.

Conversely, the study found that the impact of **Rules and Regulations (RR)** had no substantial effect on Cloud computing adoption in Saudi private hospitals, indicated by a Standardised Regression Weight (β) of 0.069 and a P-value of 0.167. This finding aligns with most previous research, particularly studies [13], [15] and underscores the complexities of technology adoption in healthcare. Factors like rules and regulations may have a different impact than initially anticipated. The insignificant impact of RR could be due to a lack

of awareness or understanding of related government rules and regulations among hospital administrators and decision-makers. If the implications of these regulations are not fully comprehended, their perceived relevance in decision-making processes related to Cloud technology adoption could be diminished. Therefore, enhancing awareness and understanding of these regulatory frameworks is critical for accurately assessing and leveraging their influence on Cloud computing adoption in the Saudi private healthcare sector.

The in-depth analysis of environmental factors finalises the study of variables affecting Cloud adoption in Saudi private hospitals, setting the stage for an updated ACC-PH framework that incorporates these insights to enhance Cloud computing deployment and use within the sector.

D. The Revised Version of The Comprehensive Framework for Adopting Cloud Computing in Private Hospitals (ACC-PH)

Based on the previous findings, the ACC-PH framework for adopting Cloud computing in Saudi private hospitals has been revised, as illustrated in Fig. 2. This updated framework integrates key insights and factors identified in our research, providing a comprehensive guide for decision-makers in these hospitals. It aims to facilitate a more effective and informed implementation of Cloud computing, addressing the specific needs and challenges of the Saudi private healthcare sector. The revised ACC-PH framework will serve as a valuable tool for hospital administrators and IT professionals, aiding them in navigating the complexities of Cloud computing adoption and ensuring its successful integration into their operations.

VI. CONCLUSION

Our study thoroughly analyses the factors influencing the adoption of Cloud computing in Saudi private hospitals, exploring technological, organisational, and environmental dimensions. Through surveying 650 managers and administrative staff, we identified several factors that significantly influence Cloud adoption, ranked by their impact: Attitude toward Change (AT), Internet Connection (IC), Competitive Pressure (CM), Security (SE), Top Management Support (TM), Cloud Providers (CP), Compatibility (CO), Prior Experience (PE), Organisational Readiness (OR), Reliability (RL), Cost Analysis (CA), and Relative Advantage (RA). Conversely, Rules and Regulations (RR) did not significantly affect adoption decisions.

The current study holds considerable importance, offering significant contributions, both theoretical and practical. Theoretically, it is groundbreaking as it explores for the first time the roles of factors such as Cloud Providers (CP) and Internet Connection (IC) within the Saudi healthcare sector. It introduces novel insights by highlighting the significant impacts of these and other investigated factors. Additionally, the study contributes confirmatory and contrasting insights on Cloud adoption theories in Saudi healthcare, enriching current discourse in this field.

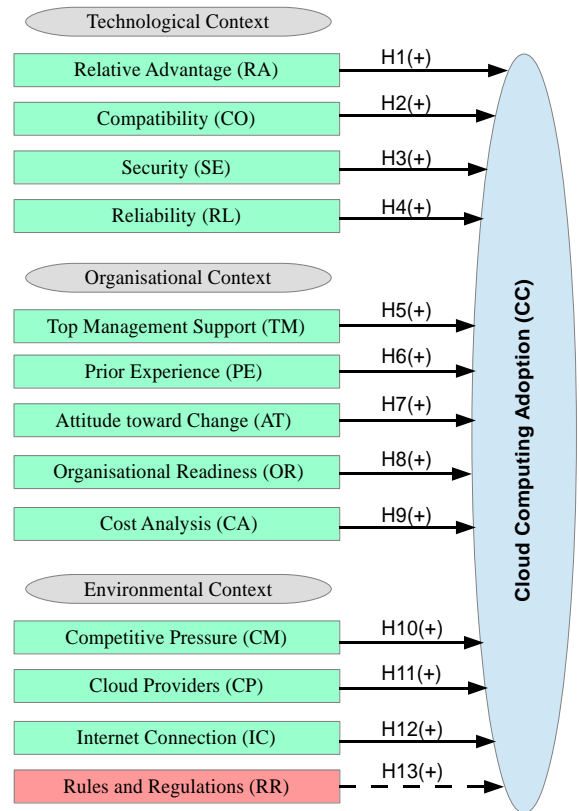


Fig. 2 The Revised Version of The Comprehensive Framework for Adopting Cloud Computing in Private Hospitals (ACC-PH)

Practically, the study is further distinguished by its exclusive focus on Saudi private hospitals, a setting previously underexplored in existing literature. This research serves as an essential guide for decision-makers in Saudi private hospitals by outlining strategic approaches for adopting Cloud computing. Implementing these strategies addresses private hospitals' dual challenges of maintaining high-quality healthcare services and enhancing cost efficiency, which are crucial for increasing profits in a competitive market. By clearly delineating the factors influencing Cloud adoption, the study equips Saudi private hospital administrators with the necessary insights to implement Cloud solutions effectively. This guidance is designed to streamline operations and enhance service delivery, leveraging technological advancements to meet operational and financial goals.

Looking ahead, we propose to enrich this research by collecting qualitative data from decision-makers at Saudi private hospitals. This further investigation will deepen our understanding of the factors influencing Cloud adoption and enhance the decision-making process. Such qualitative research will complement this study's findings and provide detailed insights to guide Cloud computing implementation strategies more effectively in the Saudi private healthcare sector.

REFERENCES

- [1] MOH Statistical Yearbook, "Statistical Yearbook," 2021. [Online]. Available:

- https://www.moh.gov.sa/en/Ministry/Statistics/book/Documents/Statistical-Yearbook-2021.pdf
- [2] R. Sajjad and M. O. Qureshi, "An assessment of the healthcare services in the Kingdom of Saudi Arabia: An analysis of the old, current, and future systems," *Int J Healthc Manag*, vol. 13, no. sup1, pp. 109–117, Dec. 2020, doi: 10.1080/20479700.2018.1433459.
 - [3] J. M. L. P. Caldeira, J. J. P. C. Rodrigues, and P. Lorenz, "Toward ubiquitous mobility solutions for body sensor networks on healthcare," *IEEE Communications Magazine*, vol. 50, no. 5, pp. 108–115, May 2012, doi: 10.1109/MCOM.2012.6194390.
 - [4] M. Khalid, H. Afzaal, S. Hassan, N. A. Zafar, S. Latif, and A. Rehman, "Automated UML-based Formal Model of E-Health System," in *2019 13th International Conference on Mathematics, Actuarial Science, Computer Science and Statistics (MACS)*, IEEE, Dec. 2019, pp. 1–6, doi: 10.1109/MACS48846.2019.9024830.
 - [5] F. Alharbi, A. S. Atkins, C. Stanier, and A. Atkins, "Strategic framework for cloud computing decision-making in healthcare sector in Saudi Arabia," in *The 7th International Conference on eHealth, Telemedicine, and Social Medicine*, 2015, pp. 138–144.
 - [6] T. U. Zaman *et al.*, "E-health and its Transformation of Healthcare Delivery System in Makkah, Saudi Arabia," *International Journal of Medical Research & Health Sciences*, vol. 7, no. 5, pp. 76–82, 2018.
 - [7] B. Loganayagi and S. Sujatha, "Enhanced Cloud Security by Combining Virtualization and Policy Monitoring Techniques," *Procedia Eng*, vol. 30, pp. 654–661, 2012, doi: 10.1016/j.proeng.2012.01.911.
 - [8] B. Balamurugan, P. Venkata Krishna, N. S. Kumar, and G. V. Rajyalakshmi, "An Efficient Framework for Health System Based on Hybrid Cloud with ABE-Outsourced Decryption," in *Artificial Intelligence and Evolutionary Algorithms in Engineering Systems*, vol. 325, Springer, New Delhi, 2015, pp. 41–49, doi: 10.1007/978-81-322-2135-7_6.
 - [9] R. T. Hameed, O. A. Mohamad, O. T. Hamid, and N. Tapus, "Design of e-Healthcare management system based on cloud and service oriented architecture," in *2015 E-Health and Bioengineering Conference (EHB)*, IEEE, Nov. 2015, pp. 1–4, doi: 10.1109/EHB.2015.7391393.
 - [10] T. Ermakova, B. Fabian, M. Kornacka, S. Thiebes, and A. Sunyaev, "Security and Privacy Requirements for Cloud Computing in Healthcare," *ACM Trans Manag Inf Syst*, vol. 11, no. 2, pp. 1–29, Jun. 2020, doi: 10.1145/3386160.
 - [11] H. Gangwar and H. Date, "Critical Factors of Cloud Computing Adoption in Organizations: An Empirical Study," *Global Business Review*, vol. 17, no. 4, pp. 886–904, Aug. 2016, doi: 10.1177/0972150916645692.
 - [12] S. T. Alharbi, "Users' acceptance of cloud computing in Saudi Arabia: An extension of Technology Acceptance Model," *International Journal of Cloud Applications and Computing*, vol. 2, no. 2, pp. 1–11, 2012, doi: 10.4018/ijcac.2012040101.
 - [13] S. S. Almubarak, "Factors influencing the adoption of cloud computing by Saudi university hospitals," *International Journal of Advanced Computer Science and Applications (IJACSA)*, vol. 8, no. 1, 2017, doi: 10.14569/ijacsa.2017.080107.
 - [14] M. O. Alassafi, "Success indicators for an efficient utilization of cloud computing in healthcare organizations: Saudi healthcare as case study," *Comput Methods Programs Biomed*, vol. 212, 2021, doi: 10.1016/j.cmpb.2021.106466.
 - [15] F. Ayadi, "Critical factors affecting the decision to adopt cloud computing in Saudi health care organizations," *Electronic Journal of Information Systems in Developing Countries*, vol. 88, no. 6, 2022, doi: 10.1002/isd2.12231.
 - [16] F. Alharbi, A. Atkins, and C. Stanier, "Understanding the determinants of cloud computing adoption in Saudi healthcare organisations," *Complex & Intelligent Systems*, vol. 2, no. 3, pp. 155–171, 2016, doi: 10.1007/s40747-016-0021-9.
 - [17] F. Alshahrani, N. Beloff, and M. White, "ACC-PH: a Comprehensive Framework for Adopting Cloud Computing in Private Hospitals," in *The 18th Conference on Computer Science and Intelligence Systems FedCSIS 2023*, Oct. 2023, pp. 17–26, doi: 10.15439/2023F4109.
 - [18] N. Alkhater, R. Walters, and G. Wills, "An empirical study of factors influencing cloud adoption among private sector organisations," *Telematics and Informatics*, vol. 35, no. 1, pp. 38–54, 2018, doi: 10.1016/j.tele.2017.09.017.
 - [19] A. N. Tashkandi and I. M. Al-Jabri, "Cloud computing adoption by higher education institutions in Saudi Arabia: An exploratory study," *Cluster Comput*, vol. 18, no. 4, pp. 1527–1537, Dec. 2015, doi: 10.1007/s10586-015-0490-4.
 - [20] J. E. Cushman, M. R. Kelly, M. Fusco-Rollins, and R. Faulkner, "Resource Review—Using Qualtrics Core XM for Surveying Youth," *Journal of Youth Development*, vol. 16, no. 1, pp. 161–167, Mar. 2021, doi: 10.5195/jyd.2021.886.
 - [21] N. J. Blunch, *Introduction to structural equation modelling using SPSS and AMOS*. Los Angeles, Calif.; London: SAGE, 2008.
 - [22] J. M. Morse, "'What's your favorite color?' Reporting Irrelevant Demographics in Qualitative Research," *Qual Health Res*, vol. 18, no. 3, pp. 299–300, 2008, doi: 10.1177/1049732307310995.
 - [23] R. B. Kline, *Principles and Practice of Structural Equation Modeling*, Fourth edition. New York: The Guilford Press, 2016.
 - [24] B. Tabachnick and L. Fidell, *Using Multivariate Statistics*, 7th ed. 2021.
 - [25] J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, *Multivariate Data Analysis*, 7th edition. Pearson, 2013.
 - [26] T. A. Brown, *Confirmatory Factor Analysis for Applied Research*, Second Edition. Guilford Publications, 2015.
 - [27] A. H. Segars, "Assessing the unidimensionality of measurement: a paradigm and illustration within the context of information systems research," *Omega (Westport)*, vol. 25, no. 1, pp. 107–121, Feb. 1997, doi: 10.1016/S0305-0483(96)00051-5.
 - [28] Z. Awang, W. M. A. Wan Afthanorhan, and M. A. M. Asri, "Parametric and Non Parametric Approach in Structural Equation Modeling (SEM): The Application of Bootstrapping," *Mod Appl Sci*, vol. 9, no. 9, Jul. 2015, doi: 10.5539/mas.v9n9p58.
 - [29] J. F. Hair, C. M. Ringle, and M. Sarstedt, "Partial Least Squares Structural Equation Modeling: Rigorous Applications, Better Results and Higher Acceptance," *Long Range Plann*, vol. 46, no. 1–2, pp. 1–12, Feb. 2013, doi: 10.1016/j.lrp.2013.01.001.
 - [30] L. Hu and P. M. Bentler, "Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives," *Struct Equ Modeling*, vol. 6, no. 1, pp. 1–55, Jan. 1999, doi: 10.1080/10705519909540118.
 - [31] J. F. Hair, C. M. Ringle, and M. Sarstedt, "PLS-SEM: Indeed a Silver Bullet," *Journal of Marketing Theory and Practice*, vol. 19, no. 2, pp. 139–152, Apr. 2011, doi: 10.2753/MTP1069-6679190202.
 - [32] J. F. Hair Jr., L. M. Matthews, R. L. Matthews, and M. Sarstedt, "PLS-SEM or CB-SEM: updated guidelines on which method to use," *International Journal of Multivariate Data Analysis*, vol. 1, no. 2, p. 107, 2017, doi: 10.1504/IJMDA.2017.10008574.
 - [33] C. Fornell and D. F. Larcker, "Evaluating Structural Equation Models with Unobservable Variables and Measurement Error," *Journal of Marketing Research*, vol. 18, no. 1, p. 39, Feb. 1981, doi: 10.2307/3151312.
 - [34] L. Cavaglione *et al.*, "Tight Arms Race: Overview of Current Malware Threats and Trends in Their Detection," *IEEE Access*, vol. 9, pp. 5371–5396, 2021, doi: 10.1109/ACCESS.2020.3048319.
 - [35] J. Gaskin and J. Lim, "Model Fit Measures," *Gaskination's StatWiki*, vol. 37, no. 3, pp. 814–822, 2016.
 - [36] L. S. Meyers, G. Gamst, and A. J. Guarino, *Applied Multivariate Research: Design and Interpretation*, 3rd ed. Thousand Oaks: SAGE Publications, Inc., 2016.
 - [37] B. M. Byrne, *Structural Equation Modeling With AMOS: Basic Concepts, Applications, and Programming*, 3rd ed. Routledge, 2016.
 - [38] R. C. MacCallum, M. W. Browne, and H. M. Sugawara, "Power analysis and determination of sample size for covariance structure modeling," *Psychol Methods*, vol. 1, no. 2, pp. 130–149, Jun. 1996, doi: 10.1037/1082-989X.1.2.130.
 - [39] A. H. Segars and V. Grover, "Re-Examining Perceived Ease of Use and Usefulness: A Confirmatory Factor Analysis," *MIS Quarterly*, vol. 17, no. 4, p. 517, Dec. 1993, doi: 10.2307/249590.

A Hybrid Machine Learning Model for Forest Wildfire Detection using Sounds

Robertas Damasevicius
Department of Applied Informatics
Vytautas Magnus University
Akademija, Lithuania
robertas.damasevicius@vdu.lt

Ahmad Qurthobi
Center of Real Time Computer Systems
Kaunas University of Technology
Kaunas, Lithuania
ahmad.qurthobi@ktu.edu

Rytis Maskeliunas
Faculty of Applied Mathematics
Silesian University of Technology
Gliwice, Poland
rytis.maskeliunas@polsl.pl

Abstract—Forest wildfires pose a significant threat to ecosystems, human settlements, and the global environment. Early detection is important for effective mitigation and response. This paper introduces a novel approach to forest wildfire detection by harnessing the unique sound signatures associated with wildfires. Our proposed model combines the strengths of deep learning techniques with heuristic optimization algorithms. The deep learning component focuses on recognizing the intricate patterns in the sound data, while the heuristic optimization, based on a Particle Swarm Optimization (PSO) algorithm, ensured the model's adaptability and efficiency in diverse forest environments. Preliminary results indicate that our hybrid model outperforms traditional methods and existing machine learning models in terms of accuracy, sensitivity, and specificity, demonstrating robustness against ambient forest noise, ensuring fewer false alarms.

Index Terms—Forest Wildfire Detection, Sound Recognition, Audio Processing, Deep Learning, Convolutional neural Network, Heuristic Optimization.

I. INTRODUCTION

FOREST wildfires have become one of the most pressing environmental challenges of the 21st century. With increasing global temperatures and changing climatic patterns, the frequency and intensity of these wildfires have seen a significant increase. The devastation caused by these fires is not limited to the loss of flora and fauna; they also have profound socioeconomic implications, affecting human settlements, agriculture, and contributing to global carbon emissions, making the timely detection and monitoring of forest wildfires a paramount task.

Historically, wildfire detection relied heavily on human observers, often stationed in lookout towers, to visually spot and report fires. As technology advanced, satellite imagery became a popular tool, offering a broader view of vast forested areas. Although satellites can provide valuable data, they come with their own set of challenges: cloud cover can obscure views and there can be delays in data acquisition and processing, which might not always allow real-time detection [1], [2].

Ground-based sensors, such as smoke detectors and infrared cameras, have also been deployed in certain high-risk areas. These systems, while effective in specific contexts, have limitations in terms of coverage and can sometimes be prone to false alarms due to other heat sources or smoke from non-wildfire sources [3], [4].

In recent years, the idea of using sound for environmental monitoring has gained attention in forests rich with distinct acoustic signatures of wildlife, vegetation, and natural phenomena such as wildfires [5]. Recognizing this, researchers have begun to explore the potential of sound-based detection systems as a complementary tool to existing methods [6], [7]. Wildfires, for example, create a distinct sound pattern that results from the combustion of materials and the rapid movement of air. The advantage of sound-based systems lies in their ability to continuously monitor an environment, unaffected by visual obstructions such as smoke or foliage [4]. With the advent of machine learning and advanced signal processing techniques, the ability to accurately distinguish between different forest sounds and pinpoint the onset of a wildfire has become a tangible reality [8], [9], [10].

The purpose of this study is to harness the potential of sound-based signatures, combined with advanced machine learning techniques, to improve the early detection of forest wildfires. Given the limitations of existing methods and the urgency of timely wildfire detection, our study seeks to explore a novel, efficient, and scalable solution, aiming to integrate a heuristic optimization algorithm with the deep learning model, aiming to enhance the adaptability, efficiency, and robustness of the model in varied forest environments.

This paper is structured as follows. Following this introduction, we present related work. Section III focuses on the methodology, detailing data collection, the deep learning model, and the integration of heuristic optimization. Section IV describes the experimental evaluation, while in Section V concludes the paper.

II. RELATED WORKS

The domain of wildfire detection has seen a number of research efforts, each aiming to harness the potential of various technological advances. Sound-based detection, while relatively new, has shown promise in recent years [11].

The idea of using sound as a detection mechanism is rooted in the understanding that every event, especially those involving rapid physical changes, such as wildfires, produces distinct acoustic signatures. Initial attempts at sound-based wildfire detection were rudimentary, relying on basic acoustic sensors to detect sudden increases in ambient noise levels

[12]. These systems were prone to false alarms, especially in noisy environments or during storm events. The integration of machine learning into sound-based wildfire detection marked a significant turning point. Algorithms capable of classifying complex sound patterns have been developed [13], [14], [15]. For example, Lee and Kim utilized Support Vector Machines (SVM) to classify forest sounds, achieving a notable accuracy in distinguishing wildfire sounds from other ambient noises [16]. Their work laid the foundation for more sophisticated models, emphasizing the potential of machine learning in this domain. Researchers also investigated distinguishing the unique sound signatures of wildfires from other forest noises [17]. Johnson and Rodriguez used Fourier transforms to analyze the frequency components of recorded sounds, successfully identifying the characteristic low-frequency rumblings of wildfires [18]. Although more accurate than its predecessors, this approach still faced challenges in real-time processing and scalability.

The deep learning application in environmental monitoring [19], [20] is now very popular, because of the complexity and vastness of environmental data, where traditional machine learning methods often struggle due to their need for manual feature extraction, since deep learning excels at automatically learning and extracting features from raw data, making it particularly suited for complex environmental datasets [19] or environmental conservation and management [19]. One of the most prominent applications of deep learning in this field is the analysis of satellite imagery. Convolutional Neural Networks (CNNs), known for their prowess in image recognition, have been used to detect changes in land cover, deforestation, and even soil moisture levels with greater precision than traditional methods [21]. Recurrent Neural Networks (RNNs), and, specifically, their variant long-short-term memory (LSTM) networks, have been instrumental in predicting air and water quality parameters. These networks are ideal for handling sequential data, making them suitable for time series environmental data [22]. For example, studies have used LSTM networks to predict air quality indices in urban areas, demonstrating the potential of deep learning for real-time environmental monitoring [22]. In addition, deep learning has found applications in the monitoring of wildlife and biodiversity. Automated systems equipped with deep learning algorithms have been developed to identify species from camera trap images, track animal movements, and even recognize bird songs, helping eco-conservation efforts and providing valuable information about ecological dynamics.

Heuristic optimization techniques, inspired by natural processes and phenomena, have been used to solve complex optimization problems [23], especially in domains where traditional methods might be computationally expensive or infeasible. Recent studies have used convolutional neural networks (CNN) to analyze sound spectrograms, achieving remarkable accuracy rates in wildfire sound detection, thanks to the optimization capabilities of heuristic methods [24]. These techniques have found significant applications in optimizing model parameters, selecting features, and enhancing overall

model performance [25], [26], and the ability to process and analyze large datasets and complex patterns [27], [28], [29].

III. METHODOLOGY

A. Data Preprocessing

1) *Sound Data Denoising*: For our methodology, we propose a combination of wavelet-based denoising and deep learning-based denoising. The wavelet method provides an initial denoising step, removing coarse-grained noise, while the autoencoder fine-tunes the denoising process, capturing and removing more subtle noise components.

Denoising is the process of removing unwanted noise from the audio signal, enhancing the signal-to-noise ratio, and ensuring that the primary focus remains on the sounds of interest, in this case, the sounds produced by wildfires. A Wavelet transform provides a multiresolution analysis of signals, as it is particularly suited for audio denoising. Given an audio signal $x(t)$, its continuous wavelet transform with respect to a wavelet $\psi(t)$ is expressed by:

$$W_x(a, b) = \int_{-\infty}^{\infty} x(t)\psi_{a,b}(t)dt \quad (1)$$

where $\psi_{a,b}(t)$ is the wavelet shifted by parameter b and scaled by parameter a . By transforming the audio signal into the wavelet domain, we can threshold the wavelet coefficients, effectively eliminating noise. The denoised signal $x_d(t)$ can then be obtained using the inverse wavelet transform.

Spectral subtraction is based on the principle of subtracting the estimated noise spectrum from the noisy signal spectrum. Given the power spectrum $P(f)$ of the noisy signal and the estimated noise power spectrum $N(f)$, the denoised signal power spectrum $D(f)$ is given by:

$$D(f) = |P(f) - \alpha N(f)| \quad (2)$$

where α is an over-subtraction factor, typically slightly greater than 1, to account for the potential underestimation of noise.

Autoencoders can be also trained to denoise audio data, so we also tried to exploit this feature. The noisy audio signal is passed through the encoder to produce a compressed representation, which the decoder then uses to reconstruct the denoised signal. Given an input noisy signal x and its denoised version x' , the reconstruction loss L is minimized:

$$L = \sum_{i=1}^N (x_i - x'_i)^2 \quad (3)$$

where N is the number of samples in the signal.

2) *Feature Extraction*: Feature extraction is a required step in transforming raw audio data into a structured format that can be processed effectively by machine learning models. By extracting salient features, we can capture the characteristics of the audio signal that are most relevant to wildfire detection. For our methodology, we propose extracting a combination of time-domain, frequency-domain, and time-frequency features,

as this comprehensive feature set ensures that our model captures the multifaceted nature of wildfire sounds, from transient crackling noises to sustained roaring sounds. These features will then serve as input to our deep learning model for classification.

Time-Domain Features are as follows:

1. *Root Mean Square Energy (RMSE)*, which quantifies the signal's energy and is given by:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N x_i^2} \quad (4)$$

where x_i is the amplitude of the signal at time i and N is the total number of samples.

2. *Zero Crossing Rate (ZCR)*, which measures the rate at which the signal changes sign. A high ZCR can indicate the presence of noise or rapid events, such as crackling fires.

$$\text{ZCR} = \frac{1}{N-1} \sum_{i=1}^{N-1} \mathbb{I}(x_i \cdot x_{i+1} < 0) \quad (5)$$

where \mathbb{I} is the indicator function.

Frequency-Domain Features are:

1. *Spectral Centroid*, which represents the center of mass of the spectrum and can be used to distinguish between low- and high-frequency sounds.

$$\text{Spectral Centroid} = \frac{\sum_{f=1}^F S(f) \cdot f}{\sum_{f=1}^F S(f)} \quad (6)$$

where $S(f)$ is the spectral magnitude at frequency f and F is the total number of frequency bins.

2. *Spectral Bandwidth*, which describes the width of the spectrum and is defined as:

$$\text{Spectral Bandwidth} = \sqrt{\frac{\sum_{f=1}^F (f - \text{Spectral Centroid})^2 \cdot S(f)}{\sum_{f=1}^F S(f)}} \quad (7)$$

3. *Mel-Frequency Cepstral Coefficients (MFCCs)* collectively represent the short-term power spectrum of a sound from a type of cepstral representation of the audio clip in the frequency domain.

Time-Frequency Representations are:

1. *Spectrogram* is a visual representation of the spectrum of frequencies in a sound signal as they vary over time. It can capture the temporal evolution of frequency components, which can be crucial to detect transient events such as wildfires.

2. *Wavelet Transform*, as discussed in the denoising section, can also be used for feature extraction. By analyzing the wavelet coefficients at different scales, we can capture both high-frequency events (such as crackling sounds) and low-frequency modulations (such as the roar of a fire).

B. Deep Learning Model

We propose a hybrid deep learning architecture that combines Convolutional Neural Networks (CNNs) for feature extraction from spectrograms with long-short-term memory (LSTM) to capture temporal dependencies in the audio data.

The architecture of the proposed model and its parameters are summarized in Table II.

TABLE I: Summary of the proposed deep learning model architecture.

Layer Type	Output Shape	Parameters	Activation
Input	$128 \times 128 \times 1$	-	-
Conv2D	$126 \times 126 \times 16$	160	ReLU
MaxPooling2D	$63 \times 63 \times 16$	-	-
Conv2D	$61 \times 61 \times 32$	4640	ReLU
MaxPooling2D	$30 \times 30 \times 32$	-	-
LSTM	30×64	24832	Tanh
Dense	30×128	8320	ReLU
Dense	30×64	8256	ReLU
Dense (Output)	30×2	130	Softmax

The input to the model is a spectrogram of the audio signal, which provides a representation of time and frequency. This allows the model to process both the spectral content of the sound and its temporal evolution. The initial layers of the model are convolutional layers designed to extract spatial features from the spectrogram. These layers can identify patterns such as the onset of a fire's crackling or the sustained energy in a fire's sound.

- *Layer 1*: 16 filters, kernel size of 3×3 , ReLU activation.
- *Layer 2*: 32 filters, kernel size of 3×3 , ReLU activation.

After each convolutional layer, a max-pooling layer reduces the spatial dimensions, focusing on the most salient features.

Following the convolutional layers, an LSTM (Long Short-Term Memory) layer captures the temporal dependencies in the audio data for recognizing patterns that evolve over time, such as the progression of a fire.

- *LSTM Layer*: 64 units, return sequences set to True.

After the recurrent layer, fully connected (dense) layers provide the capability to classify the extracted features into the desired categories (wildfire sound or non-wildfire sound).

- *Dense Layer 1*: 128 units, ReLU activation.
- *Dense Layer 2*: 64 units, ReLU activation.
- *Output Layer*: 2 units (corresponding to the two classes), softmax activation.

C. Heuristic Optimization

The weights of our neural network were optimized using Particle Swarm Optimization (PSO). This part of the optimization process was aimed at finding the set of weights that minimizes the error between the predicted and actual results during training. Each particle in the swarm represents a potential solution, that is, a specific set of weights for the entire network. The position and velocity of each particle correspond to the weights and the change in weights, respectively. The fitness function evaluates the performance of the network with a given set of weights on the training data.

PSO is inspired by the social behavior of flocking birds or the schooling of fish. In PSO, each solution in the search space is considered a "particle". These particles "fly" through the solution space with velocities that are dynamically adjusted based on their own experience and the experience of their neighbors.

The position update rule for each particle is given by:

$$x_i(t+1) = x_i(t) + v_i(t+1) \quad (8)$$

where $x_i(t)$ is the current position of the particle and $v_i(t+1)$ is its velocity at the next time step.

The velocity update rule is:

$$v_i(t+1) = w \cdot v_i(t) + c_1 \cdot \text{rand}() \cdot (pbest_i - x_i(t)) + c_2 \cdot \text{rand}() \cdot (gbest - x_i(t)) \quad (9)$$

where w is the inertia weight, c_1 and c_2 are cognitive and social scaling parameters, respectively, $pbest_i$ is the personal best position of the particle, and $gbest$ is the global best position among all particles.

D. Training and Validation

Given that our task is a binary classification (wildfire sound or non-wildfire sound), we employ the categorical cross-entropy loss function, defined as:

$$L = - \sum_{i=1}^N y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \quad (10)$$

where N is the number of classes, y_i is the true label, and p_i is the predicted probability for class i .

We used adam optimizer to dynamically adjust the learning rate during training, ensuring efficient and effective convergence.

To avoid overfitting, especially given the complexity of our model, we employ dropout regularization. Dropout layers are introduced after each dense layer, randomly setting a fraction of input units to 0 at each update during training time.


The data set was divided into three subsets. The training set, comprising 70% of the data, is used to train the deep learning model, where the Adam optimizer adjusts the model weights based on the input data to minimize the loss function. The 15% validation set is used for post-training by evaluating the model's performance after each epoch or batch. It facilitates hyperparameter tuning, namely through grid search (weights are optimized using the PSO), ensuring optimal settings like learning rates and batch sizes that prevent overfitting to the training data and promote generalization to new data. Lastly, the 15% test set remains unseen throughout model training and validation, providing an independent evaluation of the model performance.

After each epoch of training, the model performance was evaluated in the validation set. This provided an indication of how well the model is likely to perform on unseen data and helps in early stopping if the validation loss starts to increase, indicating potential overfitting. Hyperparameters that affect the model's learning process but are not directly optimized


by PSO, such as learning rates, batch sizes, and dropout rates were further fine-tuned using grid search, which allowed systematically exploring a predefined set of hyperparameter combinations to identify the configuration that maximizes the model's performance on our validation set.

E. Dataset Description

The Forest Wild Fire Sound dataset [30] includes sound recordings that capture the unique acoustic signatures associated with forest wildfires. This dataset is designed to support the development and testing of machine learning models for the detection of wildfires through sound analysis, utilizing audio data that represent various stages and intensities of forest fires. The samples of audio record used are presented in Figure 1.



(a) Signal Plot of Building on Fire Audio



(b) Spectrogram of Building on Fire Audio

Fig. 1: Audio Analysis of Building on Fire

Data processing included the extraction of various sound characteristics such as Mel frequency cepstral coefficients (MFCC), root mean square energy (RMSE), zero crossing rate (ZCR), spectral centroid and spectral bandwidth (Table II).

IV. EXPERIMENTAL EVALUATION

A. Evaluation Metrics

We have used accuracy, precision, recall and F1-score to evaluate the performance.

Accuracy represents the fraction of correctly predicted instances out of the total instances.

$$\text{Accuracy} = \frac{\text{Number of Correct Predictions}}{\text{Total Number of Predictions}} \quad (11)$$

TABLE II: Description of features extracted from the Forest Wild Fire Sound Dataset

Feature	Description
MFCCs	Mel-Frequency Cepstral Coefficients: Represents the short-term power spectrum of a sound
RMSE	Root Mean Square Energy: Quantifies the energy of the audio signal
ZCR	Zero Crossing Rate: Measures the rate at which the signal changes its sign
Spectral Centroid	Represents the center of mass of the spectrum
Spectral Bandwidth	Describes the width of the spectrum
Spectrogram	Visual representation of the spectrum of frequencies in a sound signal as they vary with time
Wavelet Transform	Captures both high-frequency events and low-frequency modulations in the sound signal

Precision and recall are crucial metrics, especially when classes are unbalanced.

- *Precision*: It represents the number of true positive predictions divided by the number of true positive and false positive predictions.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} \quad (12)$$

- *Recall (or Sensitivity)*: It represents the number of true positive predictions divided by the number of true positive and false negative actual instances.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (13)$$

- *F1-Score*: The F1-score is the harmonic mean of precision and recall. It is particularly useful when the class distribution is imbalanced.

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (14)$$

A confusion matrix was used to describe the performance of a classification model on a set of data for which the true values are known, as it provides a detailed breakdown of true positive, true negative, false positive, and false negative predictions.

The ROC was also used as a graphical representation of the true positive rate versus the false positive rate for various threshold values. In addition, we provide the AUC value that represents the degree or measure of separability, indicating how well the model distinguishes between the classes.

B. Model Performance

The performance metrics obtained are presented in Table III, showing that the model performs well in classifying wildfire sounds. The high metric values indicate that the model is reliable and effective. The high precision value (93.2%) indicates a low false-positive rate. Which is crucial for operational efficiency as it minimizes unnecessary responses to non-wildfire sounds. The high recall value (96.1%) ensures that most actual wildfires are detected. Which is critical for early intervention and minimizing the spread of wildfires. Model exhibited a balanced Performance. The high F1-Score (94.6%) showed a good balance between precision and recall, as the high AUC value (0.987) indicates that the model's performance is

robust across various threshold settings, making it versatile and reliable in different scenarios.

TABLE III: Performance of the hybrid deep learning model.

Metric	Value
Accuracy	94.7%
Precision	93.2%
Recall	96.1%
F1-Score	94.6%
AUC	0.987

The classification results are presented as confusion matrix in Figure 2 that shows the good wildfire detection performance with only a few misclassifications.

		Predicted	
		Non-Wildfire	Wildfire
Actual	Non-Wildfire	47	3
	Wildfire	2	48

Fig. 2: Confusion matrix of Forest Wildfire Detection Results

Figure 3 shows the ROC plot of the classification performance of the model, indicating that it performs well with an AUC of 0.987, showing that our hybrid approach has a strong discriminative ability in the detection of wildfires.




Fig. 3: ROC curve for Forest Wildfire Detection Results

C. Ablation Study

An ablation study has been performed to systematically evaluate the contribution of various features (e.g., MFCCs,

RMSE, ZCR) and model’s components (e.g., Conv2D layers, LSTM layers) to the overall performance of the wildfire detection model. This analysis helps to understand the impact and significance of each component in the overall performance of the model. The study involved creating several modified versions of the baseline model, each with specific features or components removed, and then measuring the resulting changes in performance metrics. The methodology of the ablation study involved the following steps:

- The complete model, integrating all features and components, was first evaluated to establish a performance benchmark.
- Individual features and network components were systematically removed or altered, creating several distinct ablation scenarios.
- For each modified model, key performance metrics—accuracy, precision, recall, F1-score, and AUC—were measured and compared against the baseline.

We have followed these experimental scenarios:

- 1) Baseline Model, includes all features and components.
- 2) Without MFCCs, removing MFCCs to assess the importance of spectral features.
- 3) Without RMSE, excluding RMSE to evaluate the significance of this time-domain feature.
- 4) Without ZCR, omitting ZCR to understand its impact.
- 5) Without Conv2D Layers, eliminating the convolutional layers to gauge their role in spatial feature extraction.
- 6) Without LSTM Layers, removing LSTMs to examine their contribution to capturing temporal dependencies.

The results of the ablation study are presented in Table IV. The baseline model, which integrates all features and components, achieved the highest performance in all metrics, confirming its robustness and effectiveness. The removal of MFCCs resulted in the most significant performance drop, underscoring the critical role of these spectral features in the capture of the unique sound signatures of wildfires. Without MFCCs, the model accuracy fell to 88.5%, precision to 85.0%, recall to 90.0%, F1-score to 87.4%, and AUC to 0.921. Excluding time-domain features like RMSE and ZCR also led to noticeable performance degradation, although to a lesser extent than MFCCs. The accuracy without RMSE and ZCR dropped to 91.2% and 92.0%, respectively, indicating that while these features are important, they are not as critical as the spectral features. The precision and recall also declined, highlighting that these features contribute to the model’s overall ability to accurately classify wildfire sounds amidst ambient noise. The removal of Conv2D layers caused a substantial reduction in performance, with accuracy decreasing to 85.3% and AUC to 0.900. This highlights the essential role of convolutional layers in extracting spatial features from spectrograms, which are crucial for identifying patterns indicative of wildfires. Similarly, the absence of LSTM layers resulted in a performance drop, although less severe than the removal of Conv2D layers. This suggests that while the temporal dependencies captured by the

LSTM layers are important, the spatial features extracted by the Conv2D layers play a more significant role in the overall performance of the model.

TABLE IV: Ablation Study Results for Wildfire Detection Model

Scenario	Accuracy	Precision	Recall	F1-Score	AUC
Baseline Model	94.7%	93.2%	96.1%	94.6%	0.987
W/o MFCCs	88.5%	85.0%	90.0%	87.4%	0.921
W/o RMSE	91.2%	89.1%	92.5%	90.8%	0.943
W/o ZCR	92.0%	89.5%	94.0%	91.7%	0.950
W/o Conv2D layers	85.3%	83.0%	88.0%	85.4%	0.900
W/o LSTM layers	89.7%	87.2%	91.0%	89.0%	0.928

Overall, the performance drop without MFCCs shows that they are crucial for high accuracy, precision, and AUC, highlighting their importance in capturing spectral features of wildfire sounds. Removing RMSE and ZCR also reduces performance, though not as significantly as MFCCs, indicating that these time-domain features contribute to the model’s robustness but are not as critical as spectral features. The significant performance decline without Conv2D layers shows their importance in extracting spatial features from the spectrograms. The drop in performance without LSTM layers indicates their importance in capturing temporal dependencies, though the impact is less severe than removing Conv2D layers. We believe these results show that the hybrid architecture sufficiently leverages both convolutional and recurrent layers, along with the combination of spectral and time-domain features, to achieve high accuracy in wildfire detection.

D. Comparison with other models

For a holistic evaluation, we compared our hybrid model with two baseline models: a pure convolutional neural network (CNN) [31], ResNet-based CNN with attention module [32], long-short-term memory (LSTM) [33], and Transformer-based Model [34]. As evident in Table V, CNN-based models, which excel in extracting spatial features from spectrograms, achieved an accuracy of 90.3%, an F1-score of 90.1%, and an AUC of 0.965. Similarly, LSTM models focusing on temporal dependencies achieved an accuracy of 88.7%, an F1-score of 88.5%, and an AUC of 0.952. Our hybrid approach achieved the best accuracy of 94.7%, an F1-score of 94.6%, and an AUC of 0.987. A ResNet-based CNN model reported an accuracy of 91.2%, an F1-score of 91.0%, and an AUC of 0.971. Using Transformer architectures achieved 92.5% accuracy, a 92.3% F1-score, and an AUC of 0.975.

TABLE V: Comparison of hybrid model with baseline models.

Model	Accuracy	F1-Score	AUC
Hybrid Model	94.7%	94.6%	0.987
CNN	90.3%	90.1%	0.965
LSTM	88.7%	88.5%	0.952
ResNet-based CNN	91.2%	91.0%	0.971
Transformer-based Model	92.5%	92.3%	0.975

E. Feature Importance Analysis

To analyze feature importance, we employ the permutation importance method. The idea is to permute the values of each feature and measure the decrease in the model performance. A larger decrease indicates a higher importance of the characteristic. Figure 4 visualizes the importance of various features. Spectral features, especially the Mel frequency cepstral coefficients (MFCCs), emerge as highly significant, followed by time-domain features such as root mean square energy (RMSE) and zero-crossing rate (ZCR). Our results show the efficacy of the proposed hybrid deep learning model in detecting forest wildfires from sound data.




Fig. 4: Feature importance plot for the hybrid deep learning model.

V. CONCLUSION

The ability to detect wildfires in their infancy stages can lead to faster response times, potentially saving vast expanses of forests and the biodiversity they house. The deployment of audio sensors in forests presents a cost-effective alternative to visual surveillance systems, allowing more extensive coverage and continuous monitoring.

The hybrid deep learning model, which combines convolutional and recurrent layers, demonstrated superior performance in capturing spatial and temporal features from the audio data. Integrating heuristic optimization techniques, particularly PSO, improved the model's performance by optimizing hyperparameters and weights, showcasing the potential of combining traditional optimization techniques with deep learning.

In our nearest future, enhancing the performance of our model is the key focus, expanding the model on the diversity and quality of the additional sound data being collected. We believe, that integrating sounds from a broader range of forest types, spanning different seasons, and encompassing varying intensities of wildfires, can significantly improve our models ability to generalize wild fire sounds across even more diverse environmental conditions. While our model demonstrated promising outcomes in controlled environments, rigorous real-world testing remains to be done. We plan to establish IoT sensors, developed at center of real time computing system, in Aukštaitijos parkas, Lithuania, for evaluating efficacy amidst

ambient forest noises, fluctuating weather patterns, and unforeseen environmental variables will be pivotal to validate its robustness and reliability in practical applications.

Moreover, the hybrid architecture of our model introduces inherent computational complexities. As part of our future research, optimizing the model for real-time processing becomes imperative second goal, particularly in environments constrained by computational resources such as the IoT edge nodes we use in the Lithuanian forests. This optimization will focus on streamlining algorithms, minimizing computational overhead, and exploring efficient hardware implementations. Hopefully, this will help enhancing the model's operational efficiency and scalability, ensuring its practical viability across a spectrum of wildfire monitoring and detection scenarios.

ACKNOWLEDGEMENT

This research paper has received funding from Horizon Europe Framework Programme (HORIZON), call Teaming for Excellence (HORIZON-WIDERA-2022-ACCESS-01-two-stage) - Creation of the centre of excellence in smart forestry "Forest 4.0" No. 101059985. This research has been co-funded by the European Union under the project "FOREST 4.0 - Ekscelencijos centras tvariai miško bioekonomikai vystyti" (Nr. 10-042-P-0002)

REFERENCES

- [1] C. Filizzola, R. Corrado, F. Marchese, G. Mazzeo, R. Paciello, N. Pergola, and V. Tramutoli. Rst-fires, an exportable algorithm for early-fire detection and monitoring: description, implementation, and field validation in the case of the msg-seviri sensor. *Remote Sensing of Environment*, 186:196–216, 2016.
- [2] Kathiravan Thangavel, Dario Spiller, R. Sabatini, S. Amici, S. T. Sasidharan, Haytham Fayek, and P. Marzocca. Autonomous satellite wildfire detection using hyperspectral imagery and neural networks: A case study on australian wildfire. *Remote. Sens.*, 15:720, 2023.
- [3] Sathishkumar Samiappan, L. Hathcock, G. Turnage, C. McCraine, J. Pitchford, and R. Moorhead. Remote sensing of wildfire using a small unmanned aerial system: Post-fire mapping, vegetation recovery and damage analysis in grand bay, mississippi/alabama, usa. *Drones*, 2019.
- [4] Shuo Zhang, Demin Gao, Haifeng Lin, and Quan Sun. Wildfire detection using sound spectrum analysis based on the internet of things. *Sensors*, 19(23):5093, Nov 2019.
- [5] E. Olteanu, V. Suci, S. Segarceanu, I. Petre, and A. Scheianu. Forest monitoring system through sound recognition. pages 75–80, 2018.
- [6] Y. Sahin and T. Ince. Early forest fire detection using radio-acoustic sounding system. *Sensors*, 9(3):1485–1498, 2009.
- [7] M. A. Sonkin, A. Khamukhin, A. Pogrebnoy, P. Marinov, Vassia Atanassova, O. Roeva, K. Atanassov, and A. Alexandrov. Intercriteria analysis as tool for acoustic monitoring of forest for early detection fires, 2018.
- [8] Alexandra Moutinho and Maria João Sousa. Transfer learning for wildfire identification in uav imagery. *Signal Processing*, 190, 2020.
- [9] A.A. Khamukhin and S. Bertoldo. Spectral analysis of forest fire noise for early detection using wireless sensor networks. 2016.
- [10] A.A. Khamukhin, A.Y. Demin, D.M. Sonkin, S. Bertoldo, G. Perona, and V. Kretova. An algorithm of the wildfire classification by its acoustic emission spectrum using wireless sensor networks. volume 803, 2017.
- [11] Olusola O. Abayomi-Alli, Robertas Damaševičius, Atika Qazi, Mariam Adedoyin-Olowe, and Sanjay Misra. Data augmentation and deep learning methods in sound classification: A systematic review. *Electronics*, 11(22), 2022.
- [12] John Smith et al. Early efforts in sound-based wildfire detection. *Journal of Environmental Monitoring*, 20(4):301–310, 1998.

- [13] V. Venkataramanan, G. Kavitha, M.R. Joel, and J. Lenin. Forest fire detection and temperature monitoring alert using iot and machine learning algorithm. pages 1150–1156, 2023.
- [14] G. Peruzzi, A. Pozzebon, and M. Van Der Meer. Fight fire with fire: Detecting forest fires with embedded machine learning models dealing with audio and images on low power iot devices. *Sensors*, 23(2), 2023.
- [15] S. Vignesh, G.M. Tarun, S. Nandi, M. Sriram, and P. Ashok. Forest fire detection and guiding animals to a safe area by using sensor networks and sound. pages 473–476, 2021.
- [16] Chang Lee and Jong Kim. Application of svm in classifying forest sounds for wildfire detection. *Journal of Machine Learning Applications*, 13(2):123–131, 2012.
- [17] T. Bhatt and A. Kaur. Automated forest fire prediction systems: A comprehensive review. 2021.
- [18] Eric Johnson and Maria Rodriguez. Use of fourier transforms for sound analysis in wildfire detection. *Journal of Acoustic Research*, 22(1):55–75, 2005.
- [19] Hailong Shu, Zhen Song, Huichuang Guo, Xi Chen, and Zhongdao Yao. Deep learning algorithms for air pollution forecasting: an overview of recent developments. *Atmosphere*, 12759:1275918 – 1275918–6, 2023.
- [20] Laura Fernandez and Raj Gupta. Deep learning models for analyzing sound spectrograms in wildfire detection. *International Journal of Deep Learning*, 4(3):200–215, 2019.
- [21] Petteri Neuvuori, Nathaniel G. Narra, Petri Linna, and T. Lipping. Crop yield prediction using multitemporal uav data and spatio-temporal deep learning models. *Remote. Sens.*, 12:4000, 2020.
- [22] Shengdong Du, Tianrui Li, Yan Yang, and S. Hornig. Deep air quality forecasting using hybrid deep learning framework. *IEEE Transactions on Knowledge and Data Engineering*, 33:2412–2424, 2018.
- [23] Noor Hassan Kadhim and Q. Mosa. Review optimized artificial neural network by meta-heuristic algorithm and its applications. *Journal of Al-Qadisiyah for Computer Science and Mathematics*, 2021.
- [24] Dawid Połap, M. Woźniak, and J. Mańdziuk. Meta-heuristic algorithm as feature selector for convolutional neural networks. *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 666–672, 2021.
- [25] Victor Stany Rozario and P. Sutradhar. In-depth case study on artificial neural network weights optimization using meta-heuristic and heuristic algorithmic approach. *AIUB Journal of Science and Engineering (AJSE)*, 2022.
- [26] D. Devikanniga, K. Vetrivel, and N. Badrinath. Review of meta-heuristic optimization based artificial neural networks and its applications. *Journal of Physics: Conference Series*, 1362, 2019.
- [27] Zhonghuan Tian and S. Fong. Survey of meta-heuristic algorithms for deep learning training. 2016.
- [28] A.K. Singh, S.M. Rafeek, P.S. Harikrishnan, and I. Wilson. Review of study on various forest fire detection techniques using iot and sensor networks. *Lecture Notes in Civil Engineering*, 301 LNCE:29–37, 2023.
- [29] K. Akyol. A comprehensive comparison study of traditional classifiers and deep neural networks for forest fire detection. *Cluster Computing*, 2023.
- [30] Forest Protection. Forest wild fire sound dataset, 2023. Accessed: 2024-02-04, URL: <https://www.kaggle.com/datasets/forestprotection/forest-wild-fire-sound-dataset>.
- [31] Kaustumbh Jaiswal and Dhairyaa Kalpeshbhai Patel. Sound classification using convolutional neural networks. In *2018 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, pages 81–84. IEEE, 2018.
- [32] Chao Yang, Xingli Gan, Antao Peng, and Xiaoyu Yuan. Resnet based on multi-feature attention mechanism for sound classification in noisy environments. *Sustainability*, 15(14):10762, 2023.
- [33] Ahmad Qurthobi and Rytis Maskeliūnas. Deep learning and acoustic approach for mechanical failure detection in industrial machinery. In *Journal of Physics: Conference Series*, volume 2673, page 012032. IOP publishing, 2023.
- [34] Shaokai Zhang, Yuan Gao, Jianmin Cai, Hangxiao Yang, Qijun Zhao, and Fan Pan. A novel bird sound recognition method based on multifeature fusion and a transformer encoder. *Sensors*, 23(19):8099, 2023.

Benchmarking OpenAI's APIs and other Large Language Models for Repeatable and Efficient Question Answering Across Multiple Documents

Elena Filipovska[†], Ana Mladenovska^{*†}, Merxhan Bajrami^{*†}, Jovana Dobрева^{*†},
Velislava Hillman[‡], Petre Lameski^{*†}, Eftim Zdravevski^{*†}

^{*}Faculty of Computer Science and Engineering, Ss. Cyril and Methodius University in Skopje, Macedonia

[†]Magix.AI, Skopje

[‡]Department of Media and Communications, London School of Economics and Political Science, London, UK

{elena, ana, jovana, merxhan.bajrami, petre, eftim}@magix.ai, {ana.mladenovska.1, bajrami.merdzan}@students.finki.ukim.mk,

{jovana.dobрева, petre.lameski, eftim.zdravevski}@finki.ukim.mk, v.hillman@lse.ac.uk

Abstract—The rapid growth of document volumes and complexity in various domains necessitates advanced automated methods to enhance the efficiency and accuracy of information extraction and analysis. This paper aims to evaluate the efficiency and repeatability of OpenAI's APIs and other Large Language Models (LLMs) in automating question-answering tasks across multiple documents, specifically focusing on analyzing Data Privacy Policy (DPP) documents of selected EdTech providers. We test how well these models perform on large-scale text processing tasks using the OpenAI's LLM models (GPT 3.5 Turbo, GPT 4, GPT 4o) and APIs in several frameworks: direct API calls (i.e., one-shot learning), LangChain, and Retrieval Augmented Generation (RAG) systems. We also evaluate a local deployment of quantized versions (with FAISS) of LLM models (Llama-2-13B-chat-GPTQ). Through systematic evaluation against predefined use cases and a range of metrics, including response format, execution time, and cost, our study aims to provide insights into the optimal practices for document analysis. Our findings demonstrate that using OpenAI's LLMs via API calls is a workable workaround for accelerating document analysis when using a local GPU-powered infrastructure is not a viable solution, particularly for long texts. On the other hand, the local deployment is quite valuable for maintaining the data within the private infrastructure. Our findings show that the quantized models retain substantial relevance even with fewer parameters than ChatGPT and do not impose processing restrictions on the number of tokens. This study offers insights on maximizing the use of LLMs for better efficiency and data governance in addition to confirming their usefulness in improving document analysis procedures.

Index Terms—OpenAI, LangChain, RAG, GPT, QA, LLM, Llama, Large Language Models, Multi-document, one-shot learning, few-shot learning Q&A

I. INTRODUCTION

TRADITIONAL document analysis methods often rely on manual review or simplistic keyword-based searches, leading to significant inefficiencies and limitations. As the volume and diversity of documents continue to expand, the need for innovative approaches that can streamline analysis while preserving accuracy and comprehensiveness arises. In this study, we explore different methods for assisting the

analysis of selected EdTech providers' Data Privacy Policy (DPP) documents. On the task at hand, we aim to evaluate the efficacy, consistency, benefits and limitations of various LLMs in assessing DPP documents. The importance for such an assessment is founded on the need for an automated, scalable and reliable way to systematically analyze large bodies of text semantically. We also aim to examine the most optimal way of using LLMs regarding the optimizing factor - whether it is the price, the duration or the format of the answers provided that plays a key role in a technical chore. All of these LLM models are trained on extensive datasets and exhibit remarkable proficiency in generating human-like responses and cognitive reasoning across diverse tasks. LLMs are built to handle various tasks, such as text generation, translation, content summary, chatbot conversations, and more [1].

Our methodology begins with the gathering of documents from a vast repository of online accessible terms of services of popular educational platforms and products. The experimental flow combining expert knowledge and automated processing is shown in Fig. 1. The first step of selecting the EdTech platforms is performed by educational experts. Likewise, the legal team defined 45 questions about key aspects of the GDPR (General Data Protection Regulation). Even though the legal expertise is crucial for other aspects of the methodology, the focus of this article is the automated processing of the documents and questions by LLM models, as shown in the top-right rectangle. The whole methodology is part of a larger study that involves manual expert validation of the AI-generated answers and evaluation of their quality.

Through a scraping and filtering process using the BeautifulSoup library, we refine this collection of edtech products, ultimately obtaining a list of about 800 DPP documents for further processing. These documents are evaluated to ensure their content aligns with our study's focus. For each DPP document, these GDPR-related questions are asked through various LLM methods, aiming to evaluate whether the documents address regulatory requirements.

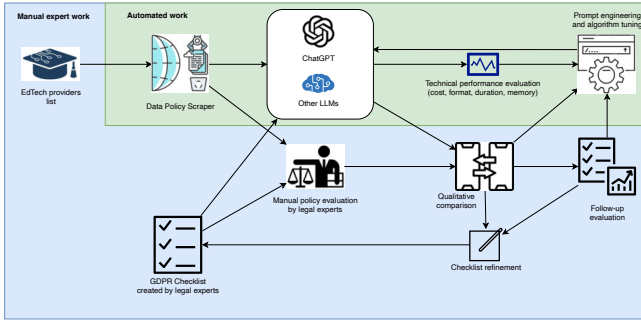


Fig. 1. Methodology combining expert knowledge and automated processing

Furthermore, we explore four distinct approaches for leveraging LLMs to analyze these documents and answer the 45 GDPR-related questions. Specifically, we explore the capabilities of OpenAI’s ChatGPT (Generative Pre-trained Transformers) GPT 3.5 Turbo, GPT 4, and GPT 4o. To use them, we evaluate OpenAI’s direct APIs and Microsoft Azure’s APIs through direct API calls, LangChain, and Retrieval Augmented Generation (RAG) systems. We also evaluate a local deployment of quantized versions (with FAISS - Facebook AI Similarity Search) of LLM models (Llama-2-13B-chat-GPTQ). Each approach offers unique advantages in processing and interpreting documents. All of these methods are systematically evaluated on multiple predefined use cases that differ in the number of documents being processed at once, the number of questions being evaluated at once, the length of the request, the length and format of the response, etc. To score the approaches, we employ a range of metrics, including response format, execution time, and cost.

Our study aims to provide insights into the optimal practices for document analysis using Artificial Intelligence (AI) technologies. By focusing on the documents themselves, we aim to uncover nuanced insights that inform decision-making and optimize outcomes in document governance and compliance. That being said, the central point is on testing AI technology for scaling otherwise manual and highly nuanced specialized literature that typically takes a long time to assess, rather than encouraging or suggesting that AI technology alone can justify if documents present companies’ compliance, transparency, or accountability to any rules or conditions in which they are mandated to operate.

This paper is structured as follows. After this introduction, Section II reviews related works, providing context and highlighting the contributions of other researchers in the field. Section III presents the methodology of the paper. Specifically, in subsection III-A we describe the data used in this study, detailing the selection and preparation processes. In subsections III-B and III-C, we outline the specific use cases designed to test the efficacy of different approaches and the evaluation methods, respectively. Subsection III-D discusses the use and implementation of OpenAI models, elaborating on each approach’s technical aspects, and Subsection III-E describes the implementation of the quantized Llama-2-13B-

chat model with FAISS vector database. Section IV presents and discusses the results, offering a detailed analysis of the findings, explains the challenges faced and the limitations of the current study. Finally, Section V concludes the paper by summarizing the key findings and suggesting directions for future research.

II. RELATED WORK

Due to the widespread use of LLMs in general domains and real world applications, various text analysis tasks have been automated and optimized, even for relatively specific tasks. Given the extensive training and massive datasets that serve as the basis for the models offered by OpenAI, their usage covers a wide set of areas. The utilization process undergoes several techniques used across almost all applications. Namely, the reading comprehension abilities of LLMs are one of their primary strengths. Ranging from classifying titles and abstracts against an inclusion/exclusion criteria, as explored in [2], to policy advising in governmental processes as elaborated in [3], the basis of implementing a system containing GPT models is text understanding.

In [4], a system is proposed that aims to address the gap in answering search queries, namely those that lack any previous internal related information about the topic. Their main approach in accomplishing this is to introduce both GPT-3.5 and GPT-4 to novel prompts, as well as instruct the model to provide a step-by-step explanation of the answering process to enhance the responses, reduce the hallucination risks, and ultimately increase the response comprehensiveness.

Authors in [5] present a question-answering chatbot aiming to answer questions related to numerous policies and guidelines set by their organization to dictate the appropriate outfit. Using the National Defence Policies and Standards of Canada as training data to their model, they developed a retrieval-based chatbot – since the responses it produces need to only be related to the aforementioned policies. The encodings of this data are based on the BERT model, so when the question vector is computed, the most similar passage of the entire corpus is retrieved, which represents the base for the answer. This response initiation by the model based on data given as additional information in the prompt is a common practice in modern AI chatbot applications and serves as the ground mechanism for Retrieval Augmented (RAG) Systems. This emerging branch of AI implies heavy reliance on prompt engineering and appropriate frameworks for optimal solutions, as well as a selection of the most relevant model for the problem at hand. Similar to this work, one of the approaches that we evaluate is the use of RAG systems to enhance the LLM responses.

Authors of [6] explore the differences in the abilities and limitations of three GPT models in the case of multiple choice question answering based on three levels of Python courses. The rapid advancements, as can be seen from the obtained results from their experiments, can significantly impact the use of generative models in the educational process. Their qualitative analysis of the results suggests very noticeable

improvements in the latest GPT model releases regarding the validity of the generated responses as well as the manner in which these responses are created.

The performance of both GPT-3.5 Turbo and GPT-4 models concerning a specific medical domain is explored in [7]. Namely, they explore the usage of the GPT models in the standardized orthopedic examination administered by the U.S orthopedic residency programs without prior exposure to similar queries. Their results indicate superior performance of the GPT-4 model over the GPT-3.5 Turbo regarding the quality of the reasoning capabilities produced.

Rapid development of AI branches always introduces security and data privacy concerns. The study [8] explores the evolving security, privacy, and data protection challenges posed by the increasing use of sophisticated digital products and platforms. It discusses how systems like chatbots process and store personal data, the security risks associated with their deployment, and the need for stringent data protection measures in line with regulations like the GDPR. Often, end-users (be that individuals or organisations) rely on data privacy policies and other lengthy texts to be reassured about how digital technologies meet data privacy and security standards and requirements. However, these texts are typically difficult to digest and comprehend and largely ignored as a result. This inspires the need to create scalable solutions that can digest these large texts and present a quick feedback response to users on whether the digital technologies they use comply and meet minimum appropriate standards and requirements. Furthermore, these concerns are one of the primary drivers in developing solutions that can be deployed in on premises or on the virtual private cloud (VPC). Therefore, one of the approaches that we evaluate is by using an LLM deployed on a local infrastructure powered by an Nvidia Titan V GPU.

III. METHODS

A. Data Description

The dataset contains data privacy policies of educational products or platforms in a textual format that describes the personal data used or collected upon registration or through the utilization of cookies, disclosure of personal information and personal data privacy, protection guidelines, and other sections that articulate how companies address legislative requirements such as the ones outlined by the GDPR. Initially, we considered over 800 products, selected based on their portrayal of personal data usage and relevance to educational establishments or online platforms [9]. Initially, each policy was scraped using the BeautifulSoup library [10]. The primary selection was made based on the policy content and type, specifically including only those policies with a minimum length of 55 characters and a non-PDF header type. This reduced the dataset to 794 policies. While evaluating the different use cases defined later in the document, we used some of these policies based on their length.

Additionally, there were defined 45 questions related to the compliance of each policy with GDPR (for full background of the work relating to the prompts development and human

analysis, see [9]). Each question was usually one sentence, such as "Is the purpose of processing the data identified?", "If you process special categories of data, do you make it clear?" etc. The list of questions and EdTech, as well as the source code for the experiments, are available at <https://github.com/admin-magix/edtech-policies/>.

The final dataset we created as a result of this research consisted of question and answer pairs received from the chatbot. Specifically, our generated dataset is a JSON format of:

```
{'Question': 'Each GDPR question',
'Answer': 'Generated answer from the chatbot'}
```

These pairs needed to be further evaluated and checked for their full accuracy.

B. Use Case Definition

To precisely evaluate the efficacy of the approaches described in Section III-D and Section III-E, we have identified six separate use cases. These vary in complexity from the simplest to the most advanced, requiring more computing power and resources. Every use case is created with the intention of thoroughly testing the strengths and weaknesses of every model inside the particular parameters of the strategy being evaluated. As a basis for each use case that we defined, we used the maximal number of tokens that the GPT 3.5-Turbo model has for a single prompt. That is to say, the primary model that we used specified the tests used for further evaluation against our factors, with each one centering around the count of policies and the count of questions implemented in one run. The following is a breakdown of the defined use cases:

- 1) **Use Case 1:** *Using one document (i.e., policy) that does not exceed the limit for the primary model token count (16K tokens) and one question.* This is relevant for shorter policies and when questions are asked one-by-one (e.g., in a scenario when a human expert is going through a checklist).
- 2) **Use Case 2:** *Using one policy that exceeds the limit for the primary model token count (16K tokens) and one question.* This is relevant for similar situations like the previous use case, but the text length of the document is much longer and can not be processed at once.
- 3) **Use Case 3:** *Using 3 policies that do not exceed the limit for the primary model token count (16K tokens) and all 45 questions.*
- 4) **Use Case 4:** *Using 10 policies that do not exceed the limit for the primary model token count (16K tokens) and 10 questions*
- 5) **Use Case 5:** *Using 50 policies that exceed the limit for the primary model token count (16K tokens) and all 45 questions*
- 6) **Use Case 6:** *Using all the policies and all the questions*

The relevance of these use cases becomes clear when evaluating models against specific metrics. Each model is analyzed based on the content of the prompts and the number

of such prompts it processes. These use cases are particularly applicable in real-life scenarios involving AI technology for analyzing large and complex texts.

C. Evaluation metrics

The defined use cases are evaluated with a list of quantitative and qualitative metrics:

- format of response given by the model (Is the response in JSON format?)
- format of the answer within the response (Are the answers firstly distinct Yes/No)
- format of the extract within the answer (If the answer is yes, is the description well extracted?)
- verification of the JSON format's correctness within the answer
- execution time
- total cost.

The Results section details how all of these metrics are assessed against a specific use case. The policies and questions follow the same format throughout all use cases in a given approach.

D. OpenAI Models

There are numerous services offered by OpenAI, most of which are based on Large Language Models (LLMs). These LLMs are deep learning models that have been trained on a very large amount of data, and as such have the ability to generate human-like responses and offer state-of-the-art cognitive reasoning on several tasks. The backbone of all the services offered relies on a set of pre-trained models. In addition to end-user services, OpenAI also provides an API - giving the opportunity to integrate the aforementioned services inside existing systems. This API acts as a link between the models offered by OpenAI and external projects, thus enabling their usage without the necessity of developing models from scratch. Users can access a range of pre-trained models, such as GPT-4o, GPT-4, GPT-3.5, DALL-E, all of which have been fine-tuned to perform specific tasks. While these models can be used as is, there is also the possibility for customization, enabling their adaptation for specific needs and fitting individual requirements. Even though projects that use these API can grow in complexity, the scalability of the infrastructure offers easy adaptation [11]. One such case is the usage of a specific pre-trained model with the aim of document analysis in comparison to a set of given questions. Access to a pre-trained model is granted after getting an OpenAI API key, which is then utilized further through four distinct approaches. Namely, the first one is a direct API call - meaning that for a specific task that needs to be accomplished a connection to OpenAI is opened, a request is made and a response is given back by the model used. This method offers a straightforward and direct way of accessing the capabilities offered by the pre-trained models. The second one focuses on using the LangChain framework, giving us a more robust way of communicating with the OpenAI models - a seamless interaction and enhanced functionality. The third one



Fig. 2. Implementation flow for a Direct API call

is implemented using a RAG system - defining a refined flow of the whole process. Ultimately, the last approach is using the API through Microsoft Azure services, aiming to simplify the integration within an existing system. The effectiveness and quality of the results given by each approach are examined by a set of metrics. By evaluating the different methods, an informed decision for the usage of the optimal practice can be made, thus enabling further project optimization.

1) *One-shot learning with direct OpenAI API call:* Using the OpenAI library and the API key created for a specific subscription on the OpenAI platform, a communication channel is instantiated, indicating the connection of our project to the pre-trained models. Once a client is established, a request is initiated for each document. Both the documents and the questions that are being answered follow a predefined format consisting of two dictionaries, where the key is the id of the document or the id of the question, and the value is the content of the document or the question, respectively. Through the process of iteration, a prompt is outlined for each of the documents - including a fixed string defining the format of the prompt and a list of the questions that need to be answered based on the current document. Considering the imposed charges for each API request, answer caching is implemented to optimize resource utilization, ensuring that only unanswered questions are included in subsequent requests. Specifically, during each iteration only those questions for the current document that have not been already answered are part of the final prompt that is given to the model. Each request includes both a system and a user message. The system message specifies the task that needs to be completed by the model - in this case it tells the model that a question-answering assignment is taking place, in which the document is described by its id and its content, and the response for each of the specified questions needs to be in a comprehensible JSON format, having a "yes" or a "no" answer and an extract denoting the part of the document that mostly contributes to the "yes" answer. Defined messages are then passed to the completions module of the client, which in turn, provides the model response (Fig. 2).

Prompt format for Direct OpenAI call and Azure OpenAI, where `questions_dict` denotes the dictionary of questions to be answered:

```

""" You are an auditor that needs to review multiple
privacy policies, each one identified with a `
PolicyID`.
For each policy, I am going to provide you the `
PolicyID` and `PolicyText`. For each policy,
answer each question that I give you with a "Yes
"/"No" answer and if the answer is "Yes" then
provide an extract from the policy that is
LONGER than 200 characters and best fits the
answer, otherwise return an empty string,

```

```

nothing else that differs from this.
Make sure that your response is only in a JSON
format like this and DO NOT PROVIDE ANY
ADDITIONAL TEXT: '{"QuestionID": {"Answer": "
Your answer", "Extract": "Extract from the
policy"}}', where " Your Answer" represents your
answer to the question, "Extract from the
policy" is the best fit extract and "QuestionID"
is the id of the question that I provide you
with (make sure it's in double quotes).
Below is the list of questions, in the following
format '{{'QuestionID': 'QuestionText'}}':{
questions_dict}
"""

```

2) *OpenAI with LangChain*: LangChain is a framework encompassing libraries and templates for developing applications powered by language models. Improving the basic connection to the OpenAI models, LangChain builds upon it by introducing the core concept of chains, representing the different AI components for creating advanced use cases based on LLMs. A chain denotes the automated actions taking place starting from the user prompt up until a response is rendered by the model. It may consist of different components, with the most frequently used ones including the prompt templates, LLMs, user agents, and the memory component [12]. In the case of document analysis, chat history was retained with the help of different prompt templates used within a memory element, which calls upon a pre-trained OpenAI model. LangChain provides several methods of recalling past interactions within a single request, and the ones tested here are `ConversationBufferMemory`, `ConversationBufferWindowMemory`, `ConversationSummaryMemory` and `ConversationSummaryBufferMemory`. All of these are built upon a `ConversationChain`, which loads context from memory. Using `ConversationBufferMemory`, we have the most basic type of conversational memory implemented - it passes the raw form of past conversation to the history parameter, which allows for subsequent prompt passing to the model. `ConversationBufferWindowMemory` offers further improvement of the conversational history by adding a window to the memory - retaining a limited number of past interactions. The `ConversationSummaryMemory` is usually used to preserve the number of tokens used in the request, by summarizing the conversation history before passing it to the history parameter. `ConversationSummaryBufferMemory` acts as a mix up of the last two types of memories providing a summarization of earlier conversation interactions while keeping the last ones in tact. After a type of memory is instantiated in the `ConversationChain`, a context for the prompt is defined, using the appropriate LangChain components in the following order:

- 1) A system message using `SystemPromptMessage`, representing the topic of the conversation between the human and the AI model
- 2) A human history message using

`HumanMessagePromptTemplate` showing a snippet of what the interaction taking place looks like so far

- 3) An AI response message using `AIMessagePromptTemplate` portraying the response given by the model
- 4) A human prompt using `HumanMessagePromptTemplate` representing the current prompt containing the policy and questions we send to the model

After a `ChatPromptTemplate` is instantiated using the list of previous messages it is sent to the conversation chain as an input, triggering the model, which gives back a response.

Human message prompt format for OpenAI with LangChain, where `questions_for_answering` denotes the dictionary of questions to be answered:

```

"""For the following policy: {policy}
answer the following questions with a "Yes
"/"No" answer and if the answer is "Yes"
then provide an extract from the policy
that is
LONGER than 200 characters and best fits
the answer, otherwise return an empty
string, nothing else that differs from
this. Make sure that your response is only
in a JSON format like this and DO NOT
PROVIDE ANY ADDITIONAL TEXT: '{{"
QuestionID": {"Answer": "Your answer", "
Extract": "Extract from the policy"}}}',
where " Your Answer" represents your
answer to the question, "Extract from the
policy" is the best fit extract and '
QuestionID' is the id of the question that
I provide you with.
Below is the list of questions, in the
following format '{{'QuestionID': '
QuestionText'}}': {questions_for_answering
}""";

```

3) *OpenAI with RAG*: The previous approaches relied heavily on the already fixed knowledge the model has been trained on. Using an approach such as a RAG system gives us the opportunity to retrieve information from an external source, further improving the reliability of the response given back by the model. Utilizing the benefits of a RAG framework indicates a combination of a retrieval component and a generational model. During the retrieval phase the system considers and fetches the most relevant piece of information given a set of sources, signifying the external information. After the most adequate piece of information is taken, it is concatenated to the input as part of the context [13]. In the specific use case of document analysis, a document represents the source from which the model answers the questions. Precisely, a single document needs to be chunked into overlapping pieces of text, which are then used during the retrieval process to find the text most adequate for the user prompt. To be able to index across a vast set of textual data, RAG depends on the concept of embedding separate chunks from the source and storing them into a vector database. Such an instance is the open-source Chroma vector database used for storing embeddings

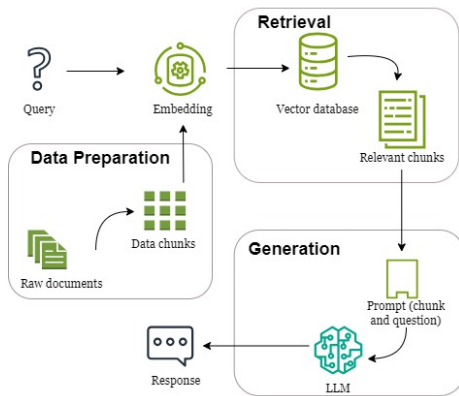


Fig. 3. Implementation flow for an OpenAI model with RAG

[14]. The flow of the whole process, starting from making a user prompt to getting a model response has several steps. Firstly, all of the questions are transformed into a vector representation, used later for indexing the chunks. Through the process of iteration a collection in the vector database is created for each document, denoting the source that is needed for the requests to the model. Having all the chunks and questions in a suitable format, querying is performed for each of the questions, meaning the closest chunk by the cosine similarity is taken. The request using the completions module from the OpenAI client has context for the prompt containing the retrieved chunk and the current question (Fig. 3).

Prompt format for OpenAI with RAG where `question_content['Text']` denotes the text of the question to be answered:

```
"""You are an auditor that needs to review
multiple privacy policies, each one identified
with a 'PolicyID'. For the following policy:"""
+ f"{closest_chunk}" + """ answer the following
question with a "Yes"/"No" answer and if the
answer is "Yes" then provide an extract
from the policy that is LONGER than 200
characters and best fits the answer, otherwise
return an empty string, nothing else that
differs from this. Make sure that your response
is only in a JSON format like this and DO
NOT PROVIDE ANY ADDITIONAL TEXT: '{"Answer": "
Your answer", "Extract": "Extract from the
policy"}', where " Your Answer"
represents your answer to the question, "
Extract from the policy" is the best fit extract
(make sure it is composed of whole sentences
and also do not include any quotation marks).
Here is the question: """ f"{question_content['
Text']}";
```

4) *Azure OpenAI*: In this approach, we leverage the Microsoft Azure Service, a collaboration between Microsoft Azure and OpenAI, which offers a comprehensive platform for developing, deploying, and managing AI-powered applications. Microsoft Azure integrates Microsoft Azure OpenAI, a powerful AI service enabling users to create and deploy AI LLMs within the Microsoft Azure platform. The Azure OpenAI Service allows developers to easily and quickly

build AI models within the Azure platform, which means that applications and usages of models can be created and deployed faster and easier within the Azure Service than traditional methods. Implementations and usages of these models occur through subscription, requiring a request access to Azure OpenAIService, which is mandatory and must be sent to Microsoft Azure for additional approval. Selecting the appropriate regions during service creation is important, as not all models are available in every region. Subsequently, we utilized Azure AI Studio to deploy models, including the creation of an Azure OpenAI GPT-4-32k model from OpenAI. Upon deployment, the model is ready to be used, the primary connectivity parameters with the model consist of an API key and Azure endpoint. These parameters were previously set up in our direct OpenAI API approach, and we reused all implementations from there.

E. Local Implementation of Llama model

In our previous approaches, we encountered several significant drawbacks:

- 1) Direct API calls proved to be the most unprofitable option and are limited by token constraints.
- 2) OpenAI's offerings require a subscription, which adds to the cost burden.

In both cases (1) and (2), if the data is sensitive and we do not want it stored externally, we cannot guarantee its privacy since it is processed on third-party systems. Additionally, challenges related to choosing the appropriate type of embedding and the strategy for Retrieval-Augmented Generation (RAG). OpenAI's API presents certain operational constraints that significantly affect its usability in constrained environments.

First, accessing OpenAI's services requires API calls limited by a predefined number of tokens. This token-based system restricts the volume of text that can be processed within a given timeframe, posing a challenge for applications with high throughput needs. Additionally, maintaining an active subscription imposes a continuous financial budget. These subscriptions, essential for accessing the API, vary in cost depending on the usage level, potentially becoming prohibitively expensive for startups and individuals with limited budgets. Furthermore, the reliance on external servers for processing raises data privacy concerns, particularly for users handling sensitive information. This reliance restricts users' control over their personal data security and makes it more difficult to comply with strict data protection laws.

To address these issues, we utilized the LLaMA-2-13B-Chat model [15] in a quantized form optimized for GPU processors¹, developed by Meta AI. It stands out due to its robust capabilities and innovative features, offering significant advantages in the realm of natural language processing:

- **Scalable Architecture**: The model is designed with a scalable architecture that can handle a wide range of computational loads, making it suitable for both high-powered

¹<https://huggingface.co/TheBloke/Llama-2-13B-chat-GPTQ>

servers and more modest local machines, depending on the deployment needs.

- **Advanced Natural Language Understanding:** With 13 billion parameters, LLaMA-2-13B-Chat demonstrates advanced understanding of complex language queries, which enhances its performance in tasks such as conversation, summarization, and text completion.
- **Efficient Memory Usage:** The quantized version of the model reduces memory requirements without a significant loss in performance, allowing it to be used effectively on devices with limited RAM and GPU resources.
- **Privacy and Security:** Local deployment capability ensures that sensitive data does not leave the organizational boundary, mitigating risks associated with data breaches and non-compliance with data protection laws.
- **Cost-Effectiveness:** By reducing dependency on cloud-based services, the LLaMA-2-13B-Chat model cuts down on ongoing operational costs related to data transmission and API usage, making it economically affordable for long-term projects.
- **Customization and Flexibility:** Users can fine-tune the model according to specific needs and constraints, which is particularly valuable in specialized applications where one-size-fits-all solutions are inadequate.

This approach allows the model to be downloaded and run locally, ensuring that the data remains within the local network. The LLaMA-2-13B-Chat model, with its 13 billion parameters, typically requires approximately 24GB of RAM to load, whereas the quantized version we used demands only 8GB. Despite this reduction, the accuracy of the responses remains consistent with the original, achieving an accuracy score of 62.18.

For building the vector base, we utilized the FAISS (Facebook AI Similarity Search) library [16], which is renowned for its efficiency in indexing and retrieving the closest matches, thereby optimizing our search and response accuracy. Specifically designed for efficient similarity searches in high-dimensional data, FAISS offers substantial advantages over Chroma for specialized applications. Its algorithms are finely tuned for vector quantization and indexing, which significantly enhances performance on both CPU and GPU—key factors for tasks that demand rapid response times and the capability to handle large datasets effectively. Moreover, FAISS integrates seamlessly with popular machine learning frameworks and benefits from robust community and developer support, ensuring it remains at the forefront of technological advancements. Also, it is the preferred choice for applications that involve complex vector operations within large-scale data environments.

By leveraging the quantized LLaMA-2-13B-Chat-GPTQ model and integrating it with the FAISS framework, we achieved a solution that maintains data privacy, reduces resource requirements, and enhances the accuracy and relevance of responses. This approach offers a valuable alternative for small business and individual users seeking to deploy advanced language models locally without compromising on

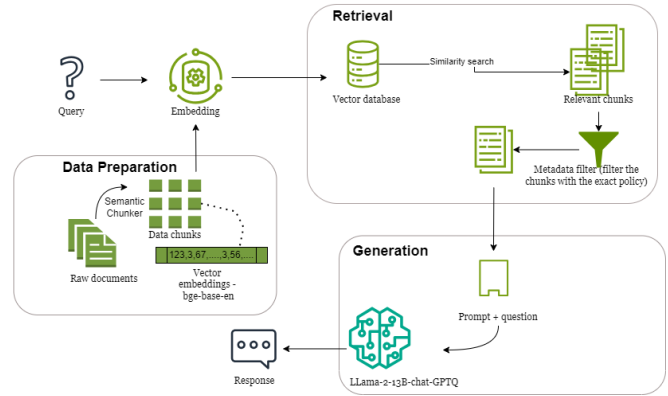


Fig. 4. Implementation flow of the Llama-2-13B-chat-GPTQ model with FAISS vector database

performance or security. The implemented architecture is given on Figure 4

Prompt format for Llama-2-13B-chat-GPTQ with FAISS:
System message for Llama-2-13B-chat-GPTQ with FAISS:

```
'''SYSTEM:You are an auditor that needs to review
multiple privacy policies, each one identified
with a 'PolicyID'. For each policy and each
question,
\\
ANSWER THE QUESTION WITH YES/NO:\\
'Yes': if you assume that the policy has the
requirements specified in the question\\
'No':otherwise\\
```

User message for Llama-2-13B-chat-GPTQ with FAISS:

```
Example:\\
USER: Policy: https://www.clarityenglish.com/privacy
.php Do you provide the information about the
identity and the contact details of the
controllers and, where applicable, of the
controller's representative? Companies which do
not have their seat in the EEA should appoint a
representative within the EU.\\
ASSISTANT:\\
```

'''

IV. RESULTS

A. Model size comparison

This section presents the results from the analysis of the previously defined Use Cases where we comparatively compare the responses of different models and methods. First, to provide a context for the evaluated models, in Table I we show the parameters, such as model architecture, number of parameters, training data and training duration.

The uniformity between the different use cases derives from the reality that a certain use case utilizes the same policies and same questions across different approaches and models. Considering this, the first use case was tested using all five approaches - two of them using only the GPT3.5-turbo with 16K tokens model, one of them using

TABLE I
COMPARISON OF EVALUATED MODELS

Feature	GPT-3.5	GPT-4	GPT-4o	Llama-213B-GPTQ
Model Architecture	Transformer	Transformer	Transformer	Transformer
Parameters (Billion)	175	280	280	213
Training Data	Diverse Text	Diverse Text	Diverse Text	Diverse Text
Training Duration	Several months	Several months	Several months	Several months
Performance	High	Very High	Optimized High	Comparable to GPT-3.5
Use Cases	NLP Tasks	Advanced NLP	Optimized NLP	NLP Tasks
Vendor	OpenAI	OpenAI	OpenAI	Meta

both the GPT3.5-turbo model with 16k tokens, the GPT4 model with 128K tokens, also the GPT4 model with 32K tokens, and at the end Llama-2-13B-chat-GPTQ with a FAISS vector database. Although the aim of the experiments with all of the evaluated methods, the content and format of the responses differs. During the evaluation phase, the GPT-4o model was utilized for the policy regulation questions. The responses provided with this model were in compliance with the desired output. The first three use cases were subject to evaluation with the GPT-4o model. Adequate responses given by the GPT-4 model were the reason for not testing the GPT-4o further, particularly because the cost for these use cases is significantly higher.

In the remainder of this section, we first evaluate the response format from an explainability point of view. Then, we directly the different models and finally, we illustrate and discuss the challenges and limitations of using OpenAI models.

B. Explainability of responses of different models

In Table II we show exemplary responses with the focus on extraction of the document content that contributed to a “Yes/No” answer for use case 1. To reiterate, use case 1 is relevant when one document (i.e., policy) does not exceed the limit for the primary model token count (16K tokens) and one question. This is relevant for shorter policies and when questions are asked one-by-one in scenarios when a human expert is going through a checklist, so the model’s output can be used to aid the expert.

Table II shows the models’ responses on following GDPR question: “Do you provide the information about the identity and the contact details of the controllers and, where applicable, of the controller’s representative? Companies which do not have their seat in the EEA should appoint a representative within the EU.”, on the ClarityEnglish policy². While all of the methods using a GPT-3.5 model return a “Yes” answer, the use case using a GPT-4 model returns a “No” answer. As evident, for the responses containing a “Yes” answer, even though there is a difference in the extract, they are all relevant to the policy for the specified question. Even though, a detailed analysis of the correctness of the results from a legal perspective and how much would they differ from the evaluation of legal experts is not in the scope of this article, we can see that different models provide different quality and length of responses. However, we point out that the

²<https://www.clarityenglish.com/privacy.php>

response of “Direct API call using GPT3.5-turbo model” and “Direct API call using GPT4 model” did not produce suitable responses.

TABLE II
EXAMPLES OF RESPONSE EXTRACT QUALITY WITH THE DIFFERENT LLM MODELS FOR USE CASE 1

Approach and model used	Use case 1 response extract
Direct API call using GPT3.5-turbo model	“Companies which do not have their seat in the EEA should appoint a representative within the EU.”
OpenAI with LangChain using GPT3.5-turbo model	“ClarityEnglish confirms that: A privacy policy is in place for the protection of personal data which has been approved by our own management and communicated to all employees and other relevant people (such as our worldwide distributors). All ClarityEnglish employees have received awareness training regarding data protection and the GDPR. Everyone within ClarityEnglish understands their roles in the protection of personal data, and will receive training where needed. We are keeping records of processing as required by the GDPR. We do not use sub-processors. All of our employees are subject to confidentiality obligations with respect to personal data. When ClarityEnglish transfers any customers’ personal data internationally, we will ensure that the transfer is legal under the GDPR. We have policies and other controls in place to provide appropriate protection of customers’ personal data, based on a careful assessment of risk.”
OpenAI with RAG using GPT3.5-turbo model	“If you have any questions, please contact our IT Manager, XYZ, who is in overall charge of privacy and data protection: XYZ@clarityenglish.com.”
Direct API call using GPT4 model	“”
Azure OpenAI using GPT4 model	“If you have any questions, please contact our IT Manager, XYZ, who is in overall charge of privacy and data protection: XYZ@clarityenglish.com. The General Data Protection Regulation (GDPR) applies to all organisations that collect and process the personal data of EU citizens.”
Llama-2-13B chat-GPTQ with FAISS	“Yes, we do provide information about the identity and contact details of our controllers and representatives as required by the GDPR. Our privacy policy includes information about our company, Clarity English Limited, including our name, address, email address, and phone number. We also identify our data protection officer (DPO) who is responsible for overseeing our compliance with data protection laws. Additionally, we inform users that they can contact us at any time to exercise their rights under the GDPR, such as accessing or deleting their personal data. We are committed to being transparent about our processing activities and to providing users with the information they need to exercise their rights under the GDPR. As a non-EEA company, we have appointed a representative within the EU, as required by Article 27 of the GDPR, to ensure that we are able to comply with the requirements of the GDPR. Our representative is [Name], who can be contacted at [Address] or [Email].”

C. Comparison of Open AI models: GPT-3.5-turbo, GPT-4, and GPT-4o

GPT-4 is recognized as a more advanced and creative release compared to the GPT-3.5-turbo model. A more detailed view into the differences between the two models based on the qualitative analysis is given in Table III.

TABLE III
COMPREHENSIVE QUALITATIVE COMPARISON AND EVALUATION OF GPT-3.5-TURBO AND GPT-4 METHODOLOGICAL APPROACHES USING SØRENSEN-DICE COEFFICIENT AND JACCARD INDEX ACROSS VARIOUS USE CASES

Approaches	Use case	Sørensen-Dice coefficient	Jaccard index
GPT-3.5 Turbo and GPT-4 Direct API call	UC 1	0	0
GPT-3.5 Turbo and GPT-4 Direct API call	UC 2	None	None
GPT-3.5 Turbo and GPT-4 Direct API call	UC 3	0.66, 0.36, 0.66	0.5, 0.22, 0.5
GPT-3.5 Turbo and GPT-4 Direct API call	UC 4	0.88, 0.9, 0.4, 0.7, 1.0, 0.2, 0.9, 0.8, 0.9, 0.9	0.66, 0.81, 0.25, 0.54, 1.0, 0.11, 0.81, 0.66, 0.81, 0.81

GPT-4o is the most advanced model and considering that the GPT-4 model provided adequate quality of responses for a cheaper cost, we did not perform significant experiments with GPT-4o. However, in the future work when legal experts will manually evaluate also the actual quality of the responses, this will be essential.

D. Comprehensive comparison of models, responses, performance and cost

Having different use cases based on the number of documents (i.e., policies) being analyzed and the number of questions being analyzed, introduces different relevant aspects for evaluation. The detailed results for each scenario tested, as can be seen from Table IV, reveal the optimal parameters for real-world circumstances. Namely, when analyzing a small body of documents and a small number of questions, per the results presented in Table IV, we conclude that the most suitable approaches are using the GPT3.5 model together with one-shot learning (i.e., direct API call) or the LangChain framework, and the Llama model. The reason for this is the sufficient quality of the responses, the appropriate response format and the low-cost of the solutions. However, as subsequent use cases were analyzed, and with that, the number of analyzed documents and the number of questions increase, it becomes apparent that the GPT4 and Llama models are more favorable in terms of response adequacy and, in some instances, the execution time. When considering the cost, all implementations with the Llama model are cheaper, assuming that the GPU-powered infrastructure is already in place.

In terms of the AI frameworks being tested, both RAG and LangChain are designed to enhance language models, but they serve different purposes and have distinct architectures. RAG, developed by Facebook AI, combines retrieval mechanisms with generation to provide more accurate and contextually relevant responses by fetching information from external documents before generating a response. This approach improves the factual accuracy and depth of generated content. LangChain, on the other hand, is a versatile framework designed to facilitate the development of applications using large language models (LLMs) by connecting them with external data sources, APIs, and custom logic, making it highly adaptable for various use cases beyond text generation, such as chatbots, information retrieval, and task automation. While RAG focuses on enhancing the quality of responses through retrieval, LangChain emphasizes the integration and utility of LLMs in diverse, complex workflows. In the evaluated use cases, regardless of the model in question, RAG show the most fitting results concerning the response quality and execution time. That being said, when evaluating these specific use cases in terms of cost, those using the Llama model are the most optimal.

E. Challenges and limitations

During the use case testing phase, we encountered numerous challenges and limitations regarding the model usage. The foremost constraint affecting the response quality is the prompt format - during the first executions, the responses of the model were either entirely not in the specified format or the content did not meet the satisfactory criteria. Across different models, the most prevalent limitation was the token limit, meaning some of the policies and questions were too long for the model to process and they returned an error. Even though these cases did not provide an answer, they still affected the

total cost. During the re-execution for some of the use cases, a difference in the answer content was noted, which proved the unpredictability of responses of LLMs. Another notable challenge when dealing with model requests is the tokens per minute limitation, which prevents seamless interaction with OpenAI.

The efficiency of `Llama-2-13B-chat-GPTQ` can be improved with a strategic selection of the vector base and embeddings. In the future, other types of embeddings could be evaluated. Additionally, implementing a graph-based RAG system, would facilitate better inference time of the responses and enable the discovery of novel relationships within the complex text data.

V. CONCLUSION

In conclusion, our study offers valuable insights into the current APIs and approaches for document analysis, for example legal documents, which is traditionally manual and time-consuming, utilizing AI technologies. Ultimately, we aim to propose a balanced approach, combining AI's capabilities with human oversight to ensure comprehensive and accurate evaluations. As a first step towards that, this paper evaluated the technical challenges related to speed, cost, and feasibility of using the current LLM technology to scale question answering across multiple documents in an automated way.

Based on the performed experiments in this study, we can conclude that using OpenAI models via API calls is a useful workaround for speeding up response times in the event that improving local infrastructure is not practical, especially for lengthy texts. On the other hand, using locally accessible quantized versions of language models, if infrastructure capacity allows, aids in data maintenance inside the local network. Furthermore, our research shows that quantized models are very relevant even if they contain less parameters than OpenAI models. Therefore, there are no restrictions on the quantity of tokens that may be processed when using quantized models.

The efficacy of using an LLM for question-answering so far are satisfactory enough to encourage further exploration into the opportunities that these models can offer. With additional optimization of model responses as well as requests to them, future uses have a potential to be more robust and even more effective. As we explore the possibilities offered by LLMs deeper, there is a growing anticipation for uncovering novel uses of integration, ultimately leading to advancements in NLP. As the boundaries of what these models can achieve expand, the promise for innovation rises. Last but not least, these outputs must be validated by human experts, so the goal is not to use human experts to simply validate AI output, but rather to use AI output to make the human work quicker and more efficient.

ACKNOWLEDGEMENTS

We would like to express our gratitude to Katarzyna Barud, Theresa Henne, Clara Saillant, Emily Thomson, and Tima Out Anwana from the University of Vienna for their exceptional

legal expertise and dedicated work in preparing the GDPR-related questions and the manual analysis of the DPP documents.

REFERENCES

- [1] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. A survey of large language models, 2023.
- [2] E Guo, M Gupta, J Deng, Y J Park, M Paget, and C Naugler. Automated paper screening for clinical reviews using large language models: Data analysis study. *J Med Internet Res*, 26, 2024.
- [3] Mehrdad Safaei and Justin Longo. The end of the policy analyst? testing the capability of artificial intelligence to generate plausible, persuasive, and useful policy analysis. *Digital Government: Research and Practice*, 5(1):1–35, 2024.
- [4] Kazem Jahanbakhsh, Mahdi Hajjibadi, Vipul Gagrani, Jennifer Louie, and Saurabh Khanwalkar. Beyond hallucination: Building a reliable question answering & explanation system with gpts.
- [5] Matt Triff Charith Gunasekara, Noah Chalifour. Question answering artificial intelligence, chatbot on military dress policy: A natural language processing based application. *Defence Research and Development Canada*, 2021.
- [6] Jaromir Savelka, Arav Agarwal, Christopher Bogart, and Majd Sakr. From gpt-3 to gpt-4: On the evolving efficacy of llms to answer multiple-choice questions for programming classes in higher education. In *International Conference on Computer Supported Education*, pages 160–182. Springer, 2023.
- [7] Michael G. Rizzo, Nathan Cai, and David Constantinescu. The performance of chatgpt on orthopaedic in-service training exams: A comparative study of the gpt-3.5 turbo and gpt-4 models in orthopaedic education. *Journal of Orthopaedics*, 2024.
- [8] Martin Hasal, Jana Nowaková, Khalifa Ahmed Saghair, Hussam Abdulla, Václav Snášel, and Lidia Ogiela. Chatbots: Security, privacy, data protection, and social aspects. *Concurrency and Computation: Practice and Experience*, 33(19):e6426, 2021.
- [9] Hillman V., Barud K., Henne T., Zdravevski E., Saillant C., and Radkoff E. In the Fine Print: Investigating EdTech Providers' Terms of Services and Data Privacy Commitment. Working Paper, 2024.
- [10] Leonard Richardson. Beautiful soup documentation, 2007.
- [11] Konstantinos I. Roumeliotis and Nikolaos D. Tselikas. Chatgpt and open-ai models: A preliminary review. *Future Internet*, 2023.
- [12] Oguzhan Topsakal and Tahir Cetin Akinci. Creating large language model applications utilizing langchain: A primer on developing llm apps fast. In *International Conference on Applied Engineering and Natural Sciences*, volume 1, pages 1050–1056, 2023.
- [13] Paulo Finardi, Leonardo Avila, Rodrigo Castaldoni, Pedro Gengo, Celio Larcher, Marcos Piau, Pablo Costa, and Vinicius Caridá. The chronicles of rag: The retriever, the chunk and the generator. *arXiv preprint arXiv:2401.07883*, 2024.
- [14] Toni Taipalus. Vector database management systems: Fundamental concepts, use-cases, and current challenges. *Cognitive Systems Research*, 2024.
- [15] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [16] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. *arXiv preprint arXiv:2401.08281*, 2024.

APPENDIX
TABLE IV
QUALITATIVE ANALYSIS OF DIFFERENT APPROACHES AND USE CASES

Model used	Approach	Use case	Response	QA result	Extract result	JSON format	Execution time	Cost
GPT Turbo	3.5 Direct call	API Use case 1	ChatCompletion object	adequate	adequate	adequate	1.7s	<\$0.01
GPT Turbo	3.5 Direct call	API Use case 2	Error code: 400 exceeded token limit	Unavailable	Unavailable	Unavailable	1.9s	<\$0.01
GPT Turbo	3.5 Direct call	API Use case 5	Error code: 400 exceeded token limit	Unavailable	partially adequate	partially adequate	5.7s	\$0.01
GPT Turbo	3.5 Direct call	API Use case 6	Error code: 400 exceeded token limit	Unavailable	partially adequate	partially adequate	2min 50s	\$0.09
GPT Turbo	3.5 OpenAI with LangChain	Use case 1	Response string	adequate	adequate	adequate	10.6s	<\$0.01
GPT Turbo	3.5 OpenAI with LangChain	Use case 5	Error code: 429	Unavailable	Unavailable	Unavailable	5.8s	<\$0.01
GPT Turbo	3.5 OpenAI with LangChain	Use case 6	Error code: 400	partially adequate	partially adequate	partially adequate	47.3s	\$0.18
GPT Turbo	3.5 OpenAI with RAG	Use case 1	Response string	adequate	adequate	adequate	1.7s	\$0.01
GPT Turbo	3.5 OpenAI with RAG	Use case 2	Response string	adequate	adequate	adequate	4.5s	\$0.01
GPT Turbo	3.5 OpenAI with RAG	Use case 3	Response string	adequate	adequate	adequate	6min 51s	\$0.16
GPT Turbo	3.5 OpenAI with RAG	Use case 4	Response string	adequate	adequate	adequate	2min 41s	\$0.14
GPT Turbo	3.5 OpenAI with RAG	Use case 5	Response string	adequate	adequate	adequate	84min 32s	\$2.86
GPT Turbo	3.5 OpenAI with RAG	Use case 6	Response string	adequate	adequate	adequate	2d 53min 2s	\$33.29
GPT 4	Direct call	API Use case 1	ChatCompletion object	adequate	adequate	adequate	4.7s	\$0.1
GPT 4	Direct call	API Use case 2	ChatCompletion object	adequate	adequate	adequate	17.9s	\$0.23
GPT 4	Direct call	API Use case 3	ChatCompletion object	adequate	adequate	adequate	5min 20s	\$0.48
GPT 4	Direct call	API Use case 4	ChatCompletion object	adequate	adequate	adequate	6min 21s	\$0.55
GPT 4	Azure OpenAI	Use case 3	Response string	adequate	adequate	adequate	44min 43.2s	\$13.2
GPT 4	Azure OpenAI	Use case 4	Response string	adequate	adequate	adequate	11min 25.3s	\$3.58
Llama-213B-GPTQ	LangChain RAG FAISS	Use case 1	Response string	adequate	adequate	adequate	5.23s	\$0
Llama-2-13B-GPTQ	LangChain RAG FAISS	Use case 2	Response string	adequate	adequate	adequate	8.21s	\$0
Llama-2-13B-GPTQ	LangChain RAG FAISS	Use case 3	Response string	adequate	adequate	adequate	12min 731s	\$0
Llama-2-13B-GPTQ	LangChain RAG FAISS	Use case 4	Response string	adequate	adequate	adequate	8min 941s	\$0
Llama-2-13B-GPTQ	LangChain RAG FAISS	Use case 5	Response string	adequate	adequate	adequate	3h 31min 27s	\$0
Llama-2-13B-GPTQ	LangChain RAG FAISS	Use case 6	Response string	adequate	adequate	adequate	2d 2h 23min 4s	\$0

Digital Twin Design for Autonomous Drones

Danish Iqbal and Barbora Buhnova

ORCID: 0000-0002-5070-5880, 0000-0003-4205-101X

Masaryk University, Faculty of Informatics

Brno, Czech Republic

Email: {danish, buhnova}@mail.muni.cz

Abstract—The rapid adoption of technology led to the rapid growth of various fields, including Unmanned Aerial Vehicles (UAV). Digital Twin (DT) became a popular concept to facilitate this progress, serving as a virtual replica of the physical drones to support run-time compliance checking, coordination, and analysis in trustworthy UAV design and operation. Nevertheless, the DT technology in UAV often lacks a precise specification and clear explanation of its characteristics, parameters, and functionalities. To address this gap, this paper investigates current research in DT applications for autonomous drones and compiles the findings towards the design of a DT to support the UAV sector. To this end, we extract the DT characteristics from existing papers and leverage these insights to propose a DT design for autonomous drones. The resulting DT is foundational in facilitating seamless collaboration and decision-making among collaborating autonomous drones in autonomous ecosystems to ensure safe and trustworthy operation, as demonstrated in a proof of concept, demonstrated through a case study of logistics shipment, showcasing the DT application for autonomous drones' collaboration in autonomous ecosystems.

Keywords: Digital Twin, Digital Twin Design, Digital Twin Properties, Autonomous Drones, Trust, Autonomous Ecosystems.

I. INTRODUCTION

DURING the past decade, the Digital Twin (DT) concept has played an important role in integrating the physical and cyber worlds of critical infrastructures [1], attracting attention in an ongoing transformation to Industry 4.0 [2], [3], and later also the smart mobility [4], [5], including its applications in the growing sector of Unmanned Aerial Vehicles (UAVs) and autonomous drones [6], [7]. In UAV applications, the DT serves as a virtual replica of a physical drone, facilitating run-time compliance checking, coordination and analysis. This capability enhances the reliability of drone operation, builds trust among collaborating drones and safe operation, and provides a transparent DT framework for decision-making. consequently, the adoption of DT technology in the UAV sector is leveraging autonomous drones for trust-building during critical operations [7].

Unfortunately, DT technology often remains vaguely described, lacking a precise specification and clear explanation of its characteristics, parameters and functionalities [8], [9], [10]. Consequently, a gap exists between theoretical conceptualizations of DTs and their practical implementations for safe communication operation among collaborating autonomous agents in autonomous ecosystems.

To address this challenge, there is a need to understand the properties, characteristics and components that build up a

DT. This involves moving beyond the generalized description towards specific methodologies and applications. Establishing clear parameters (i.e., navigation, environment perception, path planning, and control) of the DT can then support DT application in the critical operation of autonomous drones and promote the acceptance of autonomous drone ecosystems in general. Furthermore, it is foundational in facilitating seamless collaboration and decision-making among collaborating autonomous drones in autonomous ecosystems to ensure safe and trustworthy operation. DT can be a powerful tool for building trust in complex communication operations by thoroughly considering these factors.

Contribution. In this paper, we contribute to this research gap by identifying the characteristics and components of DTs for autonomous drones and proposing the process of DT design in the context of trustworthy drone collaboration in autonomous ecosystems. To this end, this work first systematically examines the application of Digital Twins (DTs) through a comprehensive review of existing literature across diverse scenarios and applications of autonomous drones and other vehicles in autonomous ecosystems. We focus on understanding how DTs act as virtual replicas and can help manage and control operations, particularly in Unmanned Aerial Vehicle (UAV) operations. We start this process with content analysis and specific search criteria to identify and analyze relevant papers on DTs applied in UAVs, particularly in autonomous drones. We then expanded our exploration to DT applications in autonomous vehicles due to the limited results identified among autonomous drones. Our aim is to gather insights from existing DT applications in autonomous vehicles to bridge existing knowledge gaps and leverage these insights to design and implement DTs for autonomous drones. The paper thoroughly examines existing literature, addressing relevant technical challenges and computational barriers reported in prior research.

Once the findings from the literature are collected, this work then analyses the findings to extract the characteristics and components of the DTs of autonomous drones, with the aim to provide a structure of the key enabling technologies towards a DT design, based on the parameters that compile a DT of an autonomous drone. As a context, we refer to a trust-building scenario where the DT is utilized in a trust-building process in drone communication [6]. Above all, the technologies used at the level of integration and the current description of DTs are a major focus of the investigation. This paper serves as an

enabler for this direction and for further work in the field of the Digital Twin in communication [2].

Paper Structure. The remainder of the paper is structured as follows. Section II gives an overview of related work, followed by the methodology of the DT characteristics extraction and design in Section III. The findings of the study results are discussed in Section IV, and the drone DT conceptual design is presented in Section V. Section VI outlines the proof-of-concept of an autonomous drone DT, and the paper concludes in Section VII.

II. RELATED WORK

Extensive research has been conducted to investigate the potential and enabler technologies of Digital Twins and to identify the key challenges encountered by practitioners during designing and implementing the Digital Twins [8], [9], [10]. Technical difficulties, computational barriers, and a shortage of well-founded frameworks and approaches have been reported as part of these challenges. The above research reviews existing literature exploring the utilization of DT design for the Unmanned Aerial Vehicle (UAV), employing Digital Twins as virtual replicas of physical objects to play a pivotal role in control and system management across various sectors, including technology, communication, and UAV operations [10].

DT in Simulation Environments. Within the state-of-the-art research, substantial attention is paid to simulation environments for DT application on shared infrastructure in future smart cities [11]. The linked approach [11] addresses safety and privacy concerns and facilitates pre-deployment testing and detection of real-time malfunctions. The research provides a simulation environment and DT support for shared drone infrastructure. It, however, lacks integration of essential functionalities for managing drone flights effectively and efficiently within airspace services. A novel DT-based intelligent collaboration framework of UAV swarms is proposed in [12]. The proposed framework establishes a high-fidelity DT model to mirror and reflect the complete life cycle of UAV swarm, integrating a machine learning algorithm to improve the decision-making and control behaviors. In [13], the authors outlined a DT-based cloud computing framework for large-scale military UAVs. The research explores factors such as business prediction, test cost, integrated perception, mission planning, and centralized control.

Data-Driven DT Concepts. A two-level data-driven DT concept for autonomous landing of aircraft is introduced in [14], featuring a DT instance for model predictive control and a real-time prototype for fluid-structure interaction and flight dynamics, though it lacks detailed construction and verification. A similar study [15] was proposed to characterize the dynamic environment of a commercial vertical take-off and landing convertible. The proposed method offers an estimation of aerodynamic forces and moments for various wind conditions but lacks realistic DT behavior due to limited functionality. Finally, a DT-based deep reinforcement learning training framework is proposed in [16] to enhance the effectiveness of a training model for UAVs by simulating complex environments

that are difficult to replicate in physical testbeds. All these approaches focus on utilizing DT in specific scenarios rather than providing a comprehensive description or detailed design of the DT itself.

DT Frameworks for UAV. Several works have explored the utilization of Digital Twins to enhance the performance of wireless communication for applications such as computation offloading, content caching, and resource sharing. Proposed DT frameworks [9] optimize UAV-based MEC servers for IoT networks, ensuring efficient task offloading with low-latency communication. In [9], a DT framework is proposed for IoT networks using UAVs for on-the-fly task offloading MEC servers in industrial automation to meet the ultra-reliable low-latency communication link requirements. The proposed DT model optimizes the communication and computation parameters, such as power and processing rates, to make UAV-URLLC task offloading efficient. Additionally, the integration of DT technology with drone-assisted data collection offers precise ship maneuvering in smart seaports, enhancing operational efficiency and reducing environmental impact. Yoon et al. [17] proposes a seismic fragility analysis using a UAV-based updated DT. Their findings showed that the proposed approach can be applied to bridge condition assessment using UAV inspection to update the DT at an abstract level. Next, Sun et al. [18] explore dynamic DT and federated learning for air-ground networks, while other studies focus on DT-driven vehicular edge computing with UAV FlexEdge [19]. A DT-driven training framework was proposed in [20]. While these studies give invaluable insight into DT frameworks for UAV, they exclusively focus on the application of DT technology, without providing details on DT design itself.

Summary. While extensive research exists that explores the application of DT as a key enabling paradigm for improving the performance of smart mobility systems, very few works can be found that reveal the actual details of the DT description and design mechanism in the context of autonomous drones. Driven by these considerations, this paper presents a design process of a DT for autonomous drones, which is grounded in a comprehensive review of the existing literature identified through our systematic search.

III. RESEARCH METHOD

This study aims to systematically explore the recent development in the application and implementation of the DT concept in autonomous drones (from 2015 onwards). The goal is to comprehensively understand the trends, advancements, and challenges in applying DT technology and designing the DT for collaborating autonomous drones. Each step of our research method is described below and illustrated in Figure 1.

A. Search Process

The first step is to identify and review related work relevant to the design of the Digital Twin relevant to the context of autonomous drones. Our search initially starts with the focus on DT in UAV and autonomous drones, with inclusion criteria

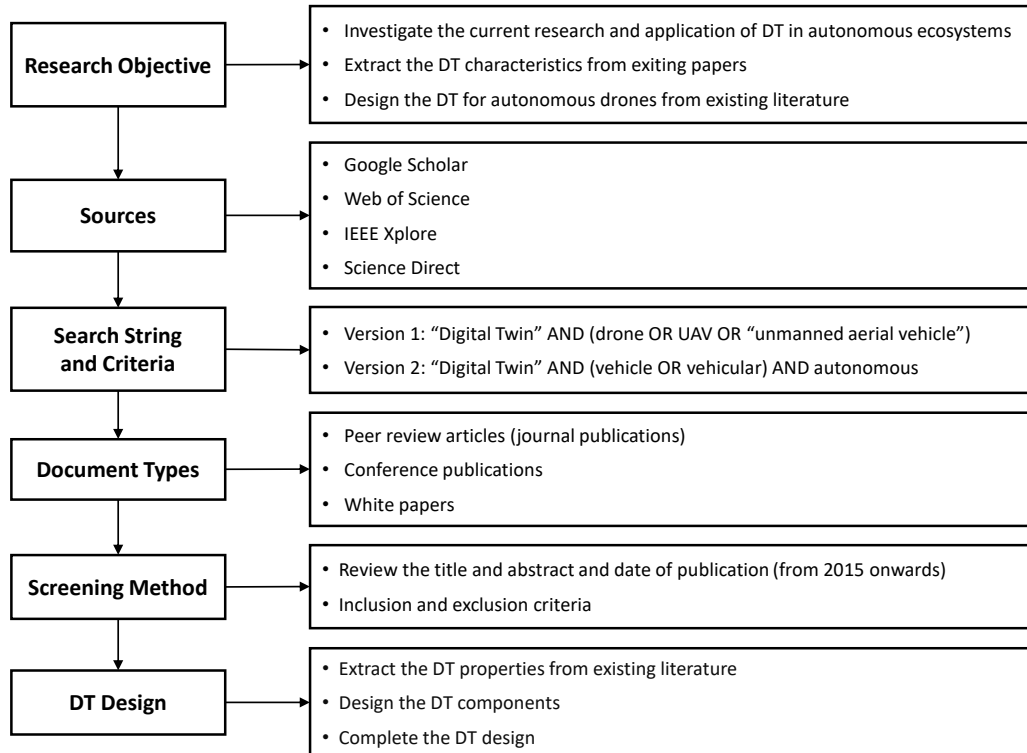


Fig. 1. Summary of Research Methodology

looking into the description of the design of a DT and its characteristics. The employed search term was:

- Version 1: “Digital Twin” AND (drone OR UAV OR “unmanned aerial vehicle”)

Due to the limited results that resulted from filtering the search results with the inclusion and exclusion criteria below, the search was expanded to its second version, exploring Digital Twin design in a wider context of autonomous vehicles, but with transferability of the findings to the narrower scope of autonomous drones. The employed search term was:

- Version 2: “Digital Twin” AND (vehicle OR vehicular) AND autonomous

We have conducted both searches across multiple academic databases, including Google Scholar, Web of Science, IEEE Explore and Science Direct, focusing on the timeframe from 2015 onwards.

B. Inclusion and Exclusion Criteria

The following inclusion and exclusion criteria were used to filter the search results.

1) Inclusion criteria (IC) – for search version 1:

- IC1: The paper provides a detailed description of a DT of a drone, including an example.
- IC2: The paper describes the structure of the DT.

2) Inclusion criteria (IC) – for search version 2:

- IC1: The paper provides a detailed description of a DT of an autonomous vehicle, including an example.
- IC2: The paper describes the structure of the DT.
- IC3: The DT is relevant for autonomous drones.

3) Exclusion criteria (EC): – for search version 1 and 2

- EC1: Papers in languages other than English.
- EC2: Gray literature (e.g., editorials and keynotes).

C. Review Execution

Although each search version started with over 100 publications (after applying the exclusion criteria and removing duplicates), the inclusion criteria reduced the search results drastically. After the inclusion criteria were applied to titles, abstracts, and full-text, only 2 papers resulted from search version 1 [9], [10], and only 7 papers resulted from search version 2 [21], [22], [23], [24], [25], [26], [27].

D. Extraction of DT Properties

Following the identification and screening of publications regarding the DT of autonomous drone applications with details description [9], [10] and the publication describing the DT details for autonomous vehicles [21], [22], [23], [24], [25], [26], [27], DT analysis was chosen to extract the DT properties to be considered in the DT design. The DT properties were extracted by analyzing text to gather information useful for collaborating drones in autonomous ecosystems. This allows

for the design and creation of patterns that autonomous drones use to exchange information during collaboration by linking predefined properties in the literature. To complement the findings resulting from this extraction, additional publications detailing autonomous drone behavior properties were considered [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38] despite not discussing them in the context of Digital Twins, to draw a more realistic and complete picture of the important properties that shall be considered in creating a virtual replica of a real autonomous drone.

E. DT Design Procedure

After extracting the properties from the literature, the design process clustered the properties to form DT components and complete the autonomous drone DT design. The workflow in the following sections details the results of this process, resulting in the DT design for autonomous drone communication in autonomous ecosystems. The DT design process is described and developed based on analyzing the selected papers to capture key concepts, trends, and hidden relationships in DT studies. The endpoint is a structured hierarchy of information and properties (the key components) of the DT for autonomous drones. The DT and its respective properties are described in Section V, based on the state-of-the-art Digital Twins properties extracted from the literature.

IV. RESEARCH STUDY RESULTS

In this section, we discuss the observations and findings extracted from the reviewed papers. Besides, we look into further observations and insights we gained when reviewing the works and describe them as opportunities for in-depth knowledge of the design of Digital Twins in the area of autonomous drones (search version 1) and autonomous vehicles (search version 2).

A. Literature Search Findings

1) *Search Version 1*: Each of the papers that passed the search and inclusion/exclusion criteria for the search version 1 is described below, describing the context in which the DT is applied and what form it takes (e.g., a state diagram, a sequence diagram, a mathematical model).

The paper [9] proposed a DT framework for Internet-of-Things (IoT) networks focusing on Unmanned Aerial Vehicles (UAVs) acting as flying Mobile Edge Computing (MEC) servers. This framework supports on-the-fly task offloading, particularly suited for industrial automation with strict constraints on Ultra-Reliable Low-Latency Communication (URLLC) links. The study formulates the end-to-end latency minimization problem for DT-aided offloading UAV-URLLC, optimizing communication and computation parameters such as power, offloading factors, and processing rates of IoT devices and MEC-UAV servers. It further details the local and edge processing aspects, including the estimation of processing rates and latency gaps between real values and DT estimations. The research approach represented the DT as a transmission

model based on a mathematical model formed by a set of equations.

The second study [10] proposed a three-layered approach to develop a comprehensive DT architecture for smart seaports, integrating physical-to-Virtual (P2V) and virtual-to-Virtual (V2V) communications. Aligned with the principle of emphasizing quality, evolution, and insight values, the architecture facilitates real-time data collection and transmission from physical entities, such as IoT devices, in the Data Layer. The Twin Layer provides digital replicas of physical entities, enabling effective monitoring and analysis of seaport operations, while the Service Layer offers insights and predictions on seaport operations, enhancing monitoring, analysis, and optimization of seaport performance. Overall, the proposed architecture enhances seaport efficiency and sustainability by providing practical insights into seaport processes and promoting sustainability in maritime logistics. The proposed study designs the DT as a state transition diagram to address the communication problem between ships and drones.

2) *Search Version 2*: Each of the papers that passed the search and inclusion/exclusion criteria for the search version 2 is described below, describing the context in which the DT is applied and what form it takes.

The study [21] presents a DT communication framework for DT-assisted Vehicular Edge Computing (VEC) networks, integrating edge computing into the vehicular network. It explores the application of DT in VEC, highlighting its benefits in real-time monitoring, adaptive network orchestration, and resource management. By integrating DT technology with Heterogeneous Information Networks (HINs), the paper presents a promising avenue for enhanced network modeling and management, ensuring resilience and adaptability to unpredictable changes in the VEC environment. Through intra-twin and inter-twin communication channels, DT-enabled VEC frameworks facilitate real-time data alignment and exchange between virtual and physical entities, thereby enhancing network performance and reliability in dynamic vehicular environments. The proposed DT is represented in the study in the form of a behavioral model that consists of network topology, channel condition and vehicle power level.

The study [22] delves into the structure and communication model of a DT, outlining its typical concept model consisting of four main components: the Physical Entity (PE), Digital Representative (DR), Intra-twin Communications, and Communications to the Outside World. In real space, the PE encompasses various infrastructures like sensors and cameras responsible for collecting real-time data of physical measurements. In the cyber domain, the DR is the application software that creates a current-model PE and takes the data it ingests and processes to render the current state and make future operation estimates. Intra-twin Communications meet the need to facilitate the interaction between the persistent execution and DR to allow the secure transport of raw data and processed information within the twin. It further discusses communications with the outside world pertaining to inter-twin communications among DRs and between DRs and the cloud.

In this way, the DR can make itself harmonic with its PEs or the sharing of information with other DRs or the cloud to maximize simulation or prediction ability and feedback loop within the DT system. As a key component of the study, the DT is represented as a digital representative model that can take part in the communication operation.

The study [23] reveals a framework of DTs in communicating vehicles based on the physical and cyber layers. The first module that is very critical in enabling communication between the two layers is using cellular technology. The physical layer incorporates vehicles, drivers, infrastructures, and sensors, while the cyber layer processes data, performs predictive analyses, and gives directives for the best maneuvers a car should take. The framework aims to improve safety, mobility, environmental sustainability and performance through integrated advanced computational techniques and seamless connectivity. The study represents the DT as a behavioral model of human motion planning and a digital replica of vehicle properties such as speed position and motion planning.

The paper [24] developed a DT framework to estimate tram position using virtual models whenever the information regarding the position is unavailable. The DT model comprises physical twin elements, which obtain sensor data from tram localization sensors and wireless routers to Digital Twins. This data is then used for virtual model training with a Long Short-Term Memory (LSTM) architecture in the Digital Twin to predict tram positions. The implementation framework includes sensor data transmission, virtual model training, data storage, and Human Machine Interface (HMI) for display integrated with the autonomous driving controller, which contains the Unscented Kalman Filter (UKF) with position estimation Stanley controller algorithm. This is done through various components embedded within the framework, such as sensors, network layers, a vehicle DT, data store, HMI, autonomous driving controller, and actuators, for its free movement and operations between physical and Digital Twins. The implementation framework of the proposed DT model is illustrated as an architecture that indicates the flow of data communication among different objects for autonomous tram localization.

The paper [25] introduces an Autonomous Vehicle Networks Digital Twin model to enable collaborative autonomous driving. Every Autonomous Vehicles (AVs) has a DT associated with an Edge Computing Device that connects it with the rest of the network. It offers a digital representation of an AV used to render support for the AV in both physical and virtual network environments. The system contains functional parameters in the AV and allows information and updates on transactions in the Autonomous Vehicular Networks (AVNs) to be made. As the AV switches between Edge Computing Devices (ECDs), its DT is passed in advance with pre-attached wired links to allow the DT to replace the AV and collaborate with other DTs on driving decisions. Collaborative driving is considered a service in the architecture of collaborative autonomous driving. The architecture is based on mapping between parameters into virtual networks of which ECDs at

the intersection contribute to determining collaborative driving decisions. Time slots are interleaved for AV coordination, and ECDs use the DTs to map across the AV parameters and co-decide on driving. The paper presents the DT in the form of a DT-enabled architecture for autonomous driving. The architecture describes the driving behavior (like lane changes, map routing, and position changes) for driving decisions.

The paper [26] introduces a DT-based trajectory prediction scheme for real-time platoon operation. The structure consists of two layers: the physical entities layer and the DT layer. Intelligent vehicles with high-precision sensors monitor the traffic in the physical entity layers. The platoon tracks surrounding social vehicles based on Long Short-Term Memory (LSTM) neural networks to predict their trajectories. The platoon members are used to distribute the responsibility to offload their load for data sensing upon an LSTM neural network used simultaneously to train the network to reduce processing delay. The DT system is operated in the head vehicle to drive the updates of the LSTM network, in which a Deep Reinforcement Learning (DRL) agent maximizes the prediction accuracy while minimizing the processing delay. The operation of the DT-based prediction scheme includes real-time data collection, trajectory prediction, DT analysis, DRL-based optimization, and execution of optimal updating strategies by the platoon. This platooning strategy aims to achieve safe and efficient driving by continuously updating the prediction model in relation to the optimized strategy and en-route data. The paper presents a behavioral DT model for autonomous vehicle platooning. They deploy the DT at the head vehicle of the platoon to reduce the processing delay in the coordination of the platoon.

The paper [27] introduced a MEC-enabled framework that embeds two kinds of domains: the physical network represented in the real world by a large domain of MEC, while a small DT sub-domain represents a virtual replica of the network. The Connected Autonomous Vehicles (CAVs) in the MEC domain are equipped with sensors and computation devices that detect line-of-sight data for making on-the-fly safety decisions. This architecture allows CAVs to intelligently orchestrate and evaluate the strategy for changing a lane to improve overall safety and efficiency. The authors proposed the DT for connected autonomous vehicles resembling a behavior model of a traffic system. It takes the form of a collection of properties such as location and road traffic prediction.

B. Extraction of DT Properties and DT Design Mapping

In this phase of the study, the comprehensive literature review on autonomous drone technology is comprehensively categorized and analyzed, including the study of the DT technology, which allows an understanding of diversified applications of their highly dynamic nature for adaptation to diverse drone technologies and safety scenarios [39], [40]. This classification process enlarges toward the application areas of the collected works in drone technology, such as design and detailed description of the DT. Finally, the extracted DT characteristics are clustered into a hierarchy of

DT components, fitting the context of DT application for trustworthy drone communication in autonomous ecosystems. The part dealing with dynamic adaptation is emphasized, examining how Digital Twins enable such adaptation in real-time by drones in environments that are diverse and constantly changing. This phase further delves into how DT defines the social metrics that enable the drone industry to adhere to the highest regulatory standards and safety metrics. This ranges from practical research into using Digital Twins in simulated and real settings to their crucial role in enhancing drone functionality and how drone technologies are continuously developed based on feedback and existing research. The results of this phase are detailed in the next section.

V. DRONE DIGITAL TWIN CONCEPTUAL DESIGN

The classification of the DT characteristics (elements, properties, and metrics) and their mapping to a hierarchy of DT components for autonomous drones used in the context of trustworthy autonomous drone communication and collaboration is presented below.

A. Properties

Each DT is characterized by a static set of properties and by the characterization of its behavior (both past and future). This section lists the static properties that form the header of the DT (as visualized in the example in Figure 4).

1) Drone ID:

- The drone ID represents a unique identifier assigned to the drone from the regulatory authority according to the aviation regulations, for instance its registration number. It is then used for identification, tracking and accountability during the flying operations.

2) Reputation:

- The reputation property refers to the trust score provided by or verifiable by a central authority of the airspace system [28].

3) Physical Condition:

- The physical condition refers to the overall status of the drone (sometimes referred to as drone health), such as:
 - Battery Power Level: The battery power level represents the status and health condition of the battery.
 - Sensor Functionality: It describes the issues in the performance and reliability of sensors.
 - Data Quality [31], [32]: Accuracy of GPS and map data for precise navigation.

B. Behavior

The core component of the drone DT is the description of its behavior, which can be specified in terms of its guarantees with various metrics as well as in terms of dynamic behavioral models (such as a Petri Nets [6], Finite State Machines FSM [7], or Functional Behavior Models FBM [36]) preferably a combination of both. This way, the DT serves as a virtual representation of the expected drone behavior. In this regard, it is noteworthy to say that such a representation

is future-directed, i.e. its examination gives clues about the declared future behavior of the autonomous drones. That is why it is useful to enrich this information in the DT with the past-behavior summary, to understand the dynamics of present-past behavioral changes of the drone, as well as past misbehavior and guarantee violations.

1) *Past Behavior Summary*: The representation and sharing of data about past behavior is the description of the combination of metrics and behavior patterns [32]. This allows autonomous drones to understand and interpret past behavior effectively.

- **Past Behavior Metrics**: Long-term trends in terms of safety, reliability, and response time, or trajectory sampling for selected past behavior are described for instance as:
 - Safety: Data on incidents, accidents, near-misses, and safety violations over a defined period.
 - Reliability: Data about past system uptime, failure rates, maintenance records, and mean time between failures (MTBF).
 - Response Time: Data on response times to incidents, customer service response times, and system performance response times.

- Safety: Data on incidents, accidents, near-misses, and safety violations over a defined period.
- Reliability: Data about past system uptime, failure rates, maintenance records, and mean time between failures (MTBF).
- Response Time: Data on response times to incidents, customer service response times, and system performance response times.

- **Past Behavior Patterns**: Behavioral patterns of both wanted and unwanted past behavior, such as:
 - Lane Change Behavior [37]: Historical data on lane change decisions, frequency, and reactions to changing traffic conditions.
 - Emergency Braking Response [34]: Documenting and analyzing instances of emergency braking, understanding triggers and outcomes.
 - Consistency Tracking [32]: Monitoring how consistently autonomous drones responded to various traffic scenarios over time.
 - Violation of aviation regulations: Summary of the violations of given guarantees or the aviation traffic laws and regulations.

2) *Declared Future Behavior*: The representation and sharing of the expected drone behavior in the present and future interactions. This allows other drones to optimize mutual interaction and also to spot suspicious behavior that might be caused by malfunction or security attacks.

- **Declared Behavior Metrics**: The drone to declares its behavior and actions for collaboration with other drones during the operation in terms of a variety of guarantees, specified with the help of relevant metrics.
 - The summary of suggested metrics is given in Section V-C.

- **Declared Behavior Patterns**: Represented in terms of a behavioral graph such as Petri nets, FSM, or FBM. These graphical models capture the sequential and concurrent aspects of the behavior, allowing autonomous drones to understand and respond to the changes in environmental conditions. That is dynamically recalibrated by the source

drone (the DT is linked to) in case of changes in its intentions. Examples include:

- Planned Trajectory [29]: Communicating intended routes, including any upcoming turns, lane changes, or stops.
- Acceleration Pattern: The drone’s expected acceleration patterns, parametrized by the surrounding objects and following the air traffic regulations.

C. Metrics

We have identified a variety of metrics relevant to autonomous drones that can be employed to support the DT design proposed in this section. The summary of the identified metrics is below. To a higher degree, these metrics are expected to detail the declared future behavior as discussed in the previous paragraph. However, the metrics can also be employed to concretize the past behavior summary.

- Quality of Service Metrics [29], [30]:
 - Dynamic Route Planning [31], [32]: The degree to which the drone plans the routes in real-time for efficiency, safety, and to avoid congestion.
 - Average velocity: The degree to which the drone follow the air traffic rules and average velocity.
 - Acceleration Smoothness [36]: The degree of smoothness of the drone acceleration under different conditions.
 - Response to Traffic Flow [36]: The degree to which the drone acceleration responds to the surrounding objects and traffic flow.
 - Emergency Response [37]: The ability to react in emergency situations, like sudden obstructions or system malfunctions is measured.
 - Sharing Data [38], [35]: The ability to share data from sensors (like LIDAR, radar, and cameras) to build a comprehensive understanding of the surrounding environment.
 - Obstacle Detection [38], [35]: The ability to communicate about detected obstacles, both static (like air traffic routes) and dynamic (like other drones and environmental risk).
- Safety Metrics [7]:
 - Distance with Front Objects: The ability to keep the distance for safety according to the rules and regulations.
 - Longitudinal Position: The ability to regulate vertical moments based on the position control (the up and down moments) according to the neighbor’s drones and other object
 - Braking Patterns [34]: The quality of the breaking in terms of the frequency, intensity and duration of braking in various scenarios.
 - Predictive Braking [34]: The ability of the drone to anticipate the need to brake based on traffic and obstacles.
 - Safety Margin Analysis [36]: The ability to maintain safe distances from other drones, considering acceleration patterns, reaction times and stopping distances.
- Regulatory Metrics [7]:
 - Lane Following: The degree of following the lane properly during the task execution operation in the airspace.
 - Speed Regulation Compliance [37]: The degree of adherence to speed limits over time and across different positions.
 - Component Interaction Compliance: The compliance of the drone with other components, such as sensors and flight controllers, which is performed according to aviation regulations.
 - Airspace Safety Regulations: Compliance with civil aviation authority’s airspace safety regulations during drone operation.
 - Licensing and Airworthiness Compliance: Compliance with civil aviation authority’s license requirements and airworthiness standards.
- Interaction/Social Metrics [33]:
 - Relative Speed to Neighbors: The degree to which the relative speed to the neighbor drone is maintained so that their safe space is protected.
 - Location [35]: The willingness to share precise location data to maintain awareness of each other’s positions.
 - Maneuver Plans [37]: The willingness to notify other drones of planned maneuvers, such as exiting air traffic or entering new zones.
 - Turn Signal [37]: The willingness to Indicate when a turn or lane change is about to occur.
 - Sharing Data [38], [35]: The willingness to share data from sensors (like LIDAR, radar, and cameras) to build a comprehensive understanding of the surrounding environment.
 - Obstacle Detection [38], [35]: The willingness to Communicate about detected obstacles, both static (like air traffic routes) and dynamic (like other drones and environmental risk).

VI. PROOF OF CONCEPT

The proof of concept is formulated in terms of DT application in a concrete case study, relying on the employment of an autonomous drone DT for trustworthy drone communication in the scenario of UAV logistics.

A. Case Study

Technological advancements drastically changed the shopping concept as the retail industry adopted online shopping methods. This new way of shopping possesses enormous changes in the process of logistics companies and retailers to provide their goods to the customers without delays and mistakes. Thus, managing online shipment to the relevant

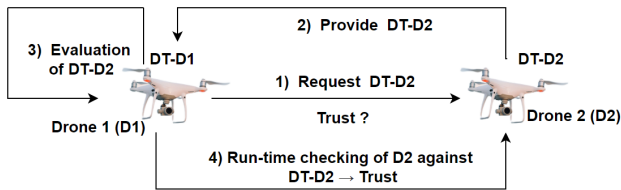


Fig. 2. How Drone 1 Decides to Trust Drone 2 [6]

customer to be operated with automated methods. In this regard, deploying an autonomous drone in logistics services requires an ecosystem where a large number of drones needs to coordinate with each other to avoid collisions, which might be caused by faulty (unintentionally harmful) or malicious (intentionally harmful) drone behaviour [6]. To address this issue, we propose utilizing the exchange of information between agents to communicate their mutual behavior, where the information is exchanged in form of their DTs.

The core idea of the approach is illustrated in Figure 2 with an interaction of two drones. The assessment of trust from the perspective of Drone 1 towards Drone 2 signifies a critical process employing a Digital Twin (DT) for evaluation. Drone 1 views Drone 2 as a black box that might be malicious. Now, Drone 1 requests Drone 2 to declare its behavior as a DT. Drone 1 checks this DT (declared behavior), and if it does not indicate any suspicious actions, it allows Drone 2 to enter its proximity while observing Drone 2 real-time behavior and its compliance with the declared behavior (communicated in form of Drone 2 DT). If Drone 2 actions align with the provided DT, trust is maintained during the operation; otherwise, trust is withdrawn. Subsequently, the incident is reported to the authority for the implementation of operational safeguards and further investigation.

B. Implementation Environment

Consider the following implementation setup, supporting the logistics scenario where the main function of the autonomous drone is to carry a shopping parcel from one place to another. As the foundation, the autonomous drone is equipped with an autonomous flight system. The autonomous drone is connected to the ground station through a communication data link and then assists with the exchange of information with other elements such as related personnel, and isolation space management [41]. Each autonomous drone utilizes the DT for the exchange of information and elements with other drones, according to the communication situation and environmental conditions. The physical space is equipped with an environment and virtual interaction platform in a real-time interaction state. In physical space, various sensors, agents, and actuators are used to rely on the perception of autonomous drones and uncertain environments and perceived information to provide feedback to a virtual data interaction center through communication. Finally, the control data interaction center sends the data to the physical terminal of the autonomous

drone to make decisions. Hence, constructing an autonomous drone DT communication channel is very important. Figure 3 displays the composition and elements of the scenario. Through the DT communication channel, the autonomous drone can be controlled, and communication and transmission can be realized.




Fig. 3. Implementation Environment

C. Digital Twin Design of an Autonomous Drone

Based on the investigation of the case study and the DT characteristics (elements, properties, and metrics), this section gives a simple example of a DT to support the given scenario, i.e. to support trustworthy autonomous drone communication and collaboration in autonomous ecosystems. Note that while the design is simple, it is not unnecessarily too simple, as also in reality, the drones are expected to exchanged only a limited set of information (in an attempt to limit the level of information disclosure towards a potentially untrusted ecosystem member, and to optimize transmission speed). The DT design is presented below, with a structured presentation in Figure 4.

1) Properties:

- Drone ID:
 - Registration Number: X95804
- Reputation:
 - Trust Score: 0.94
- Physical Condition:
 - Battery Power Level: 100%
 - Sensor Functionality: 90%

2) Behavior:

- Past Behavior Summary:
 - Past Behavior Metrics:
 - * Safety, i.e. percentage of safe operations, calculated as: $(1 - (2 \text{ incidents} + 1 \text{ accident} + 1 \text{ near miss} + 0 \text{ safety violations}) / 57 \text{ total operations}) = 93\%$.
 - * Reliability, i.e. data about past system uptime (98%), failure rate (0.5%), maintenance records,

Digital Twin Design for Autonomous Drone			
A) Properties	1) Drone ID	Registration Number: X95804	
	2) Reputation	Trust Score: 0.94	
	3) Physical Condition	PC1: Battery Power Level: 100% PC2: Sensor Functionality: 90%	
B) Behavior			
1) Past Behavior Summary	• Past Behavior Metrics	PBM1: Safety: 93%	
		PBM2: Reliability: 90%	
	• Past Behavior Patterns	PBP1: Lane Change Behavior: data with snapshots of all unsuccessful lane changes.	
		PBP2: Emergency Braking Response: documenting and analyzing instances of emergency braking, understanding triggers and outcomes.	
PBP3: Violation of Aviation Regulations: on 05/01/2024 violation of RegX125, on 07/03/2024 violation of RegX248, 08/06/2024 violation of RegX584.			
2) Declared Future Behavior	• Declared Behavior Metrics	Quality of Service Metrics	QSM1: Adaptive Route Planning: 93%
			QSM2: Acceleration Smoothness: 88%
		Interaction/Social Metrics	SIM1: Relative Speed to Neighbors: 97%
	SIM2: Location: 95%		
	SIM3: Turn Signal: 98%		
	• Declared Behavior Patterns	DBP1: Acceleration Pattern is described through FSM in Figure 5.	
		DBP2: Turn Strategy states are illustrated through FSM in Figure 5.	
		DBP3: Navigation and Mapping planning are described through FSM in Figure 5.	
DBP4: Operation Pattern is presented through FSM in Figure 5.			

Fig. 4. Proposed Digital Twin Design for the Autonomous Drone Case Study

and MTBF, weighted and aggregated to a reliability level of 90%.

- Past Behavior Patterns:
 - * Details are provided in Figure 4.
- Declared Future Behavior:
 - Declared Behavior Metrics:
 - * Adaptive Route Planning, i.e. the declared ratio of effective adaptive route planning instances in a subsequent block of 100 instances, calculated as (93 effective planning instances / 100 total instances) = 93%.
 - * Acceleration Smoothness, i.e. the declared ratio of smooth acceleration instances in a subsequent block of 100 instances, calculated as (88 smooth acceleration instances / 100 total instances) = 88%.
 - * Relative Speed to Neighbors, i.e. the declared ratio of instances when the indicated relative speed to neighbours is maintained in a subsequent block of 100 instances, calculated as (97 successful relative speed instances / 100 total instances) = 97%.
 - * Location, i.e. the declared ratio of instances where the precise location is successfully shared in a subsequent block of 100 instances, calculated as

(95 instances of successful precise data sharing / 100 total instances) = 95%.

- * Turn Signal, i.e. the declared ratio of the instances where signal changes are successfully indicated in a subsequent block of 100 instances, calculated as (98 successfully indicated instances / 100 total instances) = 98%.
- Declared Behavior Patterns:
 - * Details are provided in Figure 4, together with the behavior patterns for the proposed scenario illustrated via FSM in Figure 5.

The designed DT can further be used in different scenarios of exchanging information during collaboration in autonomous ecosystems. For instance, when each autonomous drone enters the neighboring zone, it can start an exchange of information with another drone through the DT to perform safe operations, which are described in the next section. Every autonomous drone shares a set of properties or functionalities in the form of DT to achieve certain goals and designed objectives during the collaboration in autonomous ecosystems.

D. Trust Assurance through Digital Twin Evaluation

At this point, the autonomous drones are physically deployed and controlled by the central authority, each commu-

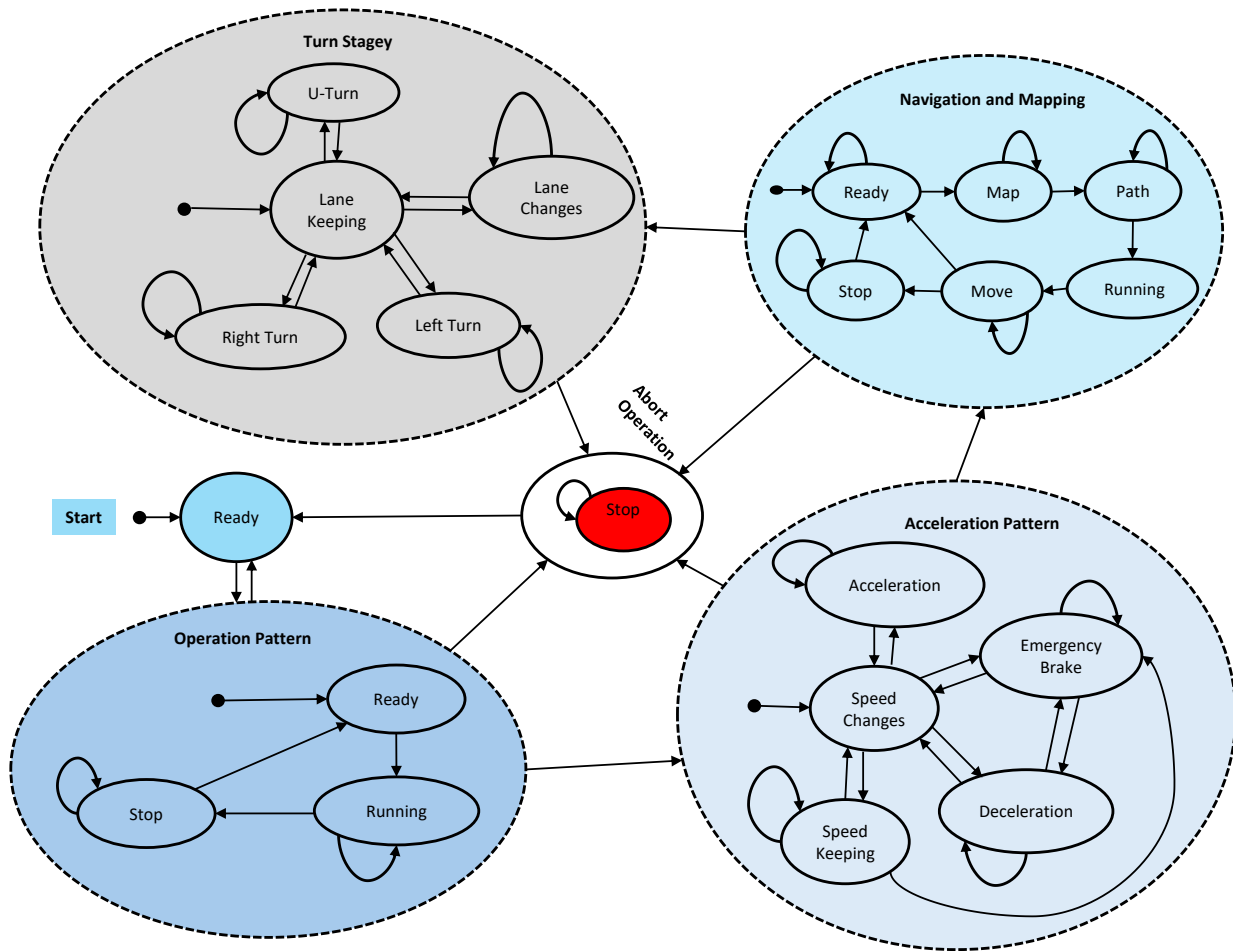


Fig. 5. Declared Behavior Patterns of an Autonomous Drone via FSM (Inspired from [42], [43])

nicating its DT as discussed above. After the deployment, the different number of drones will start to execute their own tasks. During the task execution, the collaboration is needed. During the operation, multiple drones pass each other in a space of autonomous ecosystems. To avoid collision and ensure safe operation, the autonomous drone requests its peer to declare the intended behavior for trust assessment and share it with in form of its DT.

Based on understanding their intended behavior, it can better plan the collaboration and ensure its safety. Furthermore, it can understand the deviations from the intended behavior to report the mismatch to the authorities that can take further action to protect the ecosystem from misbehaving drones. In this regard, reward and punishment strategy (upgrading or downgrading the Trust Score of the drone) can be employed in lighter cases, and direct drone isolation in the more serious ones.

VII. CONCLUSION

The primary contribution of this research is the investigation of existing literature on the DT design for autonomous drones. As also seen after the investigation of the existing literature,

the details discretion of DT for an autonomous drone is still in infancy as literature mainly focuses on the application of DT at the abstract level rather than concrete details, description, and functionalities for an autonomous drone. After studying the relevant literature on properties, elements, and functionalities, we proposed a DT design for autonomous drones. In the end, we presented the DT design along with a detailed design characterization in a proof of concept with a simple example of an autonomous drone logistics shipment to illustrate the DT in a concrete scenario. The proposed DT design attempts to serve as a foundation for autonomous drones in facilitating seamless collaboration in decision-making to ensure safe and trustworthy operations of drones in autonomous ecosystems, and offers a stepping stone for further research in the domain. In the future, we want to use the DT design along with properties for the experimental exchange of information among collaborating autonomous drones, deployed in a realistic environment simulation.

ACKNOWLEDGMENT

The work was supported by GAMU project "Forensic Support for Building Trust in Smart Software Ecosystems" (no. MUNI/G/1142/2022).

REFERENCES

- [1] A. Bierska, B. Buhnova, and H. Bangui, "An integrated checklist for architecture design of critical software systems," in *FedCSIS (Position Papers)*, 2022, pp. 133–140.
- [2] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital twin in manufacturing: A categorical literature review and classification," *Ifac-PapersOnline*, vol. 51, no. 11, pp. 1016–1022, 2018.
- [3] R. Capilla, E. Cioroai, B. Buhnova, and J. Bosch, "On autonomous dynamic software ecosystems," *IEEE Transactions on Engineering Management*, vol. 69, no. 6, pp. 3633–3647, 2021.
- [4] K. Wang, Z. Li, T. Yu, and K. Sakaguchi, "Smart mobility digital twin for automated driving: Design and proof-of-concept," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*. IEEE, 2023, pp. 1–6.
- [5] H. Bangui, B. Buhnova, D. Kusnirakova, and D. Halasz, "Trust management in social internet of things across domains," *Internet of Things*, vol. 23, p. 100833, 2023.
- [6] D. Iqbal and B. Buhnova, "Model-based approach for building trust in autonomous drones through digital twins," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022, pp. 656–662.
- [7] D. Iqbal, B. Buhnova, and E. Cioroai, "Digital twins for trust building in autonomous drones through dynamic safety evaluation," in *Proceedings of the 18th International Conference on Evaluation of Novel Approaches to Software Engineering - ENASE, INSTICC*. SciTePress, 2023, pp. 629–639.
- [8] K. Singh, B. Hazarika, C.-P. Li, K. F. Tsang, and S. Biswas, "Digital twin-assisted resource allocation in uav-aided internet of vehicles networks," in *2023 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2023, pp. 409–414.
- [9] Y. Li, D. V. Huynh, T. Do-Duy, E. Garcia-Palacios, and T. Q. Duong, "Unmanned aerial vehicle-aided edge networks with ultra-reliable low-latency communications: A digital twin approach," *IET Signal Processing*, vol. 16, no. 8, pp. 897–908, 2022.
- [10] Y. Yigit, L. D. Nguyen, M. Ozdem, O. K. Kinaci, T. Hoang, B. Canberk, and T. Q. Duong, "Twinport: 5g drone-assisted data collection with digital twin for smart seaports," *Scientific reports*, vol. 13, no. 1, p. 12310, 2023.
- [11] N. Grigoropoulos and S. Lalis, "Simulation and digital twin support for managed drone applications," in *2020 IEEE/ACM 24th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*. IEEE, 2020, pp. 1–8.
- [12] L. Lei, G. Shen, L. Zhang, and Z. Li, "Toward intelligent cooperation of uav swarms: When machine learning meets digital twin," *Ieee Network*, vol. 35, no. 1, pp. 386–392, 2020.
- [13] Y.-c. Wang, N. Zhang, H. Li, and J. Cao, "Research on digital twin framework of military large-scale uav based on cloud computing," in *Journal of Physics: Conference Series*, vol. 1738, no. 1. IOP Publishing, 2021, p. 012052.
- [14] A. McClellan, J. Lorenzetti, M. Pavone, and C. Farhat, "A physics-based digital twin for model predictive control of autonomous unmanned aerial vehicle landing," *Philosophical Transactions of the Royal Society A*, vol. 380, no. 2229, p. 20210204, 2022.
- [15] D. Aláez, X. Olaz, M. Prieto, J. Villadangos, and J. J. Astrain, "Vtol uav digital twin for take-off, hovering and landing in different wind conditions," *Simulation Modelling Practice and Theory*, vol. 123, p. 102703, 2023.
- [16] S. Li, X. Lin, J. Wu, A. K. Bashir, and R. Nawaz, "When digital twin meets deep reinforcement learning in multi-uav path planning," in *Proceedings of the 5th International ACM Mobicom Workshop on Drone Assisted Wireless Communications for 5G and Beyond*, 2022, pp. 61–66.
- [17] S. Yoon, S. Lee, S. Kye, I.-H. Kim, H.-J. Jung, and B. F. Spencer Jr, "Seismic fragility analysis of deteriorated bridge structures employing a uav inspection-based updated digital twin," *Structural and Multidisciplinary Optimization*, vol. 65, no. 12, p. 346, 2022.
- [18] W. Sun, N. Xu, L. Wang, H. Zhang, and Y. Zhang, "Dynamic digital twin and federated learning with incentives for air-ground networks," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 1, pp. 321–333, 2020.
- [19] B. Li, W. Xie, Y. Ye, L. Liu, and Z. Fei, "Flexedge: Digital twin-enabled task offloading for uav-aided vehicular edge computing," *IEEE Transactions on Vehicular Technology*, 2023.
- [20] G. Shen, L. Lei, X. Zhang, Z. Li, S. Cai, and L. Zhang, "Multi-uav cooperative search based on reinforcement learning with a digital twin driven training framework," *IEEE Transactions on Vehicular Technology*, 2023.
- [21] S. R. Jeremiah, L. T. Yang, and J. H. Park, "Digital twin-assisted resource allocation framework based on edge collaboration for vehicular edge computing," *Future Generation Computer Systems*, vol. 150, pp. 243–254, 2024.
- [22] T. H. Luan, R. Liu, L. Gao, R. Li, and H. Zhou, "The paradigm of digital twin communications," *arXiv preprint arXiv:2105.07182*, 2021.
- [23] Z. Wang, X. Liao, X. Zhao, K. Han, P. Tiwari, M. J. Barth, and G. Wu, "A digital twin paradigm: Vehicle-to-cloud based advanced driver assistance systems," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE, 2020, pp. 1–6.
- [24] Y. T. Mulyadi, M. R. R. Putra, Y. Y. Nazaruddin, and M. I. Mandasari, "Digital twin model development for autonomous tram localization," *International Journal of Sustainable Transportation Technology*, vol. 5, no. 2, pp. 55–60, 2022.
- [25] Y. Hui, X. Ma, Z. Su, N. Cheng, Z. Yin, T. H. Luan, and Y. Chen, "Collaboration as a service: Digital-twin-enabled collaborative and distributed autonomous driving," *IEEE Internet of Things Journal*, vol. 9, no. 19, pp. 18 607–18 619, 2022.
- [26] H. Du, S. Leng, J. He, and L. Zhou, "Digital twin based trajectory prediction for platoons of connected intelligent vehicles," in *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. IEEE, 2021, pp. 1–6.
- [27] B. Fan, Y. Wu, Z. He, Y. Chen, T. Q. Quek, and C.-Z. Xu, "Digital twin empowered mobile edge computing for intelligent vehicular lane-changing," *IEEE Network*, vol. 35, no. 6, pp. 194–201, 2021.
- [28] J. Guo, Q. Yang, S. Fu, R. Boyles, S. Turner, and K. Clarke, "Towards trustworthy perception information sharing on connected and autonomous vehicles," in *2020 International Conference on Connected and Autonomous Driving (MetroCAD)*. IEEE, 2020, pp. 85–90.
- [29] K. Wong, Y. Gu, and S. Kamijo, "Mapping for autonomous driving: Opportunities and challenges," *IEEE Intelligent Transportation Systems Magazine*, vol. 13, no. 1, pp. 91–106, 2020.
- [30] I. Yaqoob, L. U. Khan, S. A. Kazmi, M. Imran, N. Guizani, and C. S. Hong, "Autonomous driving cars in smart cities: Recent advances, requirements, and challenges," *IEEE Network*, vol. 34, no. 1, pp. 174–181, 2019.
- [31] A. Majd, M. Loni, G. Sahebi, and M. Daneshalab, "Improving motion safety and efficiency of intelligent autonomous swarm of drones," *Drones*, vol. 4, no. 3, p. 48, 2020.
- [32] E. Cheung, A. Bera, E. Kubin, K. Gray, and D. Manocha, "Classifying driver behaviors for autonomous vehicle navigation," 2018.
- [33] W. Wang, L. Wang, C. Zhang, C. Liu, L. Sun *et al.*, "Social interactions for autonomous driving: A review and perspectives," *Foundations and Trends® in Robotics*, vol. 10, no. 3-4, pp. 198–376, 2022.
- [34] Y. Fu, C. Li, F. R. Yu, T. H. Luan, and Y. Zhang, "A decision-making strategy for vehicle autonomous braking in emergency via deep reinforcement learning," *IEEE transactions on vehicular technology*, vol. 69, no. 6, pp. 5876–5888, 2020.
- [35] C. Dong and J. M. Dolan, "Continuous behavioral prediction in lane-change for autonomous driving cars in dynamic environments," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 3706–3711.
- [36] R. Cuer, L. Piétrac, E. Niel, S. Diallo, N. Minoiu-Enache, and C. Dang-Van-Nhan, "A formal framework for the safe design of the autonomous driving supervision," *Reliability Engineering & System Safety*, vol. 174, pp. 29–40, 2018.
- [37] E.-Y. Kang, D. Mu, L. Huang, and Q. Lan, "Verification and validation of a cyber-physical system in the automotive domain," in *2017 IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C)*. IEEE, 2017, pp. 326–333.
- [38] S. Hakak, T. R. Gadekallu, P. K. R. Maddikunta, S. P. Ramu, M. Parimala, C. De Alwis, and M. Liyanage, "Autonomous vehicles in 5g and beyond: A survey," *Vehicular Communications*, vol. 39, p. 100551, 2023.

- [39] D. Halasz and B. Buhnova, "Rethinking safety in autonomous ecosystems." in *FedCSIS (Position Papers)*, 2022, pp. 81–87.
- [40] M. Macak, S. Bojnak, and B. Buhnova, "Identification of unintentional perpetrator attack vectors using simulation games: A case study," in *2021 16th conference on computer science and intelligence systems (FedCSIS)*. IEEE, 2021, pp. 349–356.
- [41] Z. Kaleem and M. H. Rehmani, "Amateur drone monitoring: State-of-the-art architectures, key enabling technologies, and future research directions," *IEEE Wireless Communications*, vol. 25, no. 2, pp. 150–159, 2018.
- [42] S.-H. Bae, S.-H. Joo, J.-W. Pyo, J.-S. Yoon, K. Lee, and T.-Y. Kuc, "Finite state machine based vehicle system for autonomous driving in urban environments," in *2020 20th International Conference on Control, Automation and Systems (ICCAS)*. IEEE, 2020, pp. 1181–1186.
- [43] N. D. Van, M. Sualeh, D. Kim, and G.-W. Kim, "A hierarchical control system for autonomous driving towards urban challenges," *Applied Sciences*, vol. 10, no. 10, p. 3543, 2020.

Critical Success Factors for ERP Projects Revisited: An Update of Literature Reviews

Christian Leyh, Alisa Lorenz
0000-0003-0535-0336
0000-0002-8547-1391
Technical University of Central
Hesse (THM) – University of
Applied Sciences
THM Business School
Wiesenstr. 14,
35390 Gießen, Germany
Email: {christian.leyh,
alisa.lorenz}@w.thm.de

Michael Jan Faruga
Bachelor Graduate
Technical University of Central
Hesse (THM) – University of
Applied Sciences
THM Business School
Wiesenstr. 14,
35390 Gießen, Germany

Linda Koller
Bachelor Graduate
Technische Universität Dresden
Chair of Information Systems, esp.
IS in Industry and Trade
Helmholtzstr. 10,
01069 Dresden, Germany

Abstract—The aim of our study was to provide an update to the research field of critical success factors (CSFs) of enterprise resource planning (ERP) system projects. Therefore, we conducted two literature reviews, more specifically systematic reviews of relevant articles in different databases and among several international conference proceedings from mid-2012 to 2023. These two reviews serve as an update of previous analyses that we conducted covering the period from 1998 until mid-2012. In our current analysis, we identified 272 relevant papers – single or multiple case studies, surveys, and literature reviews or articles from which CSFs can be derived. From these existing studies, we discovered 33 different CSFs for ERP projects. The top three factors identified are 1) Top management support and involvement, 2) Organizational fit of the ERP system, and 3) User training. Our analysis shows that the majority of important CSFs are organizational factors with a strategic perspective. Nevertheless, important technological CSFs should also be decisively considered in ERP projects. These CSFs must be coordinated and aligned with the organizational factors in order to contribute to the holistic success of the ERP project.

Index Terms—ERP systems, ERP projects, Enterprise resource planning, Critical Success Factors, CSF, Literature Review

I. INTRODUCTION

INFORMATION and communication technology (ICT) has undergone rapid changes over the last ten years, with far-reaching effects on almost every aspect of daily life. These changes and the associated digitalization have led to companies as well as public institutions, organizations, and private individuals increasingly shifting their activities to the digital world. More and more, companies are using e-commerce to tap into new market opportunities. Public authorities are turning to e-government to reduce the administrative burden, and social interactions are increasingly shifting to the digital world of social networks. Even a large proportion of global financial transactions are almost inconceivable without the use of ICT [1].

For companies, digitalization opens up the opportunity to develop and establish new business models. Companies can make their products and services “smarter” by connecting them to digital Internet of Things (IoT) platforms and collecting and analyzing data. This approach makes it possible to create new forms of value based on cross-company data exchange, the sale of data-based expertise, or data-supported collaboration on digital platforms. The consideration of data-based value creation activities through digitalization offers companies significant competitive advantages, enables more sustainable economic activity, and creates more diverse and attractive jobs in respective regions. This development underlines the role of digitalization as a key component for growth and success in the modern, dynamic economic world [2], [3], [4], [5].

Although digitalization affects all industries at different speeds and to different degrees, hardly any company can escape this development. In order to remain competitive, companies must actively shape their own digitalization and utilize the opportunities that arise [6], [7]. Therefore, for companies, an in-depth understanding of digital transformation in general, and digital innovations as well as ICT in particular is essential. This is the only way for companies to fully utilize the potential of new technologies and develop innovative business models. The ability to recognize digital trends at an early stage and integrate them into the corporate strategy in a targeted manner is crucial to future success. Companies must, therefore, continuously invest in training their employees, optimizing their processes, and, above all, adapting their own enterprise system landscape to leverage the potential of digital transformation [8], [9], [10], [11], [12], [13].

From a more technical viewpoint, the focus is primarily on enterprise systems (ES) as company-wide application systems. Over the last two decades, both technological and organizational innovations have significantly expanded the functionalities of ES. Despite decades of experience, these

implementation, adaptation, and integration projects remain complex and challenging, as changes in the ES landscape often have a profound impact on corporate structures and processes. Well-thought-out project management, embedded in long-term digital transformation strategies, is, therefore, essential for any significant change to the ES landscape in companies and public administrations [14].

Enterprise resource planning (ERP) systems are one of the most important classes of ES. For many companies, ERP systems form the basis for optimizing internal processes and are, therefore, essential for dealing with the challenges of digitalization. The implementation or adaptation of ERP systems are complex and time-consuming projects, which offer companies great opportunities, but also harbor considerable risks. To avoid being overwhelmed by these risks, companies must focus on the factors that can influence such projects. Studies have shown that paying attention to so-called “critical success factors” (CSFs) can have a positive influence on the success of ERP projects and effectively minimize risks. Paying attention to these CSFs is also relevant for other digitalization projects. Companies should consider these in every software or digitalization project, be it the implementation of new systems and technologies or the adaptation or replacement of existing systems. Therefore, the central research question underlying our previous and current analyses of CSFs is:

Which critical success factors influence ERP implementation and adaptation projects?

Critical success factors for ERP projects have already been examined in numerous scientific and practice-oriented publications. Various case studies, surveys and the results of literature reviews have been published. In our own literature analyses [15], [16], we have previously shown that CSFs for ERP projects and their importance to these projects change over time. We have, therefore, emphasized the need to update the literature reviews at certain intervals. To address this and update the existing reviews by including current ERP literature and to provide insight into the current discussion on CSFs for ERP projects, we conducted two literature reviews. The first was done in 2017 (covering the period from mid-2012 until mid-2017), and the second review was done in 2023 (covering the period from mid-2017 until 2023). More specifically, these were two systematic reviews of articles from different databases and from several international conference proceedings. The CSFs reported in this paper were derived from 272 papers identified as relevant, and the frequency of the occurrence of each CSF was counted. The aggregated results of these reviews will be presented in this paper.

This article is structured as follows. The next section presents a short overview of our data collection methodology in order to make our review reproducible. Subsequently, in Section 3, the most important critical success factors that were focused on during the reviews will be explained in detail.

Additionally, Section 3 deals with the results of the literature reviews. We will point out which factors are the most important and which factors seem to have little influence on ERP project success. Finally, the paper concludes with a summary of the results as well as the critical acclaim of the conducted literature reviews.

II. DATA COLLECTION METHODOLOGY – SYSTEMATIC LITERATURE REVIEW

The literature reviews to identify the CSFs were performed in several steps similar to the approach suggested by Webster & Watson [17]. The steps taken in both current reviews are strongly based on the procedure that we used in our previous literature reviews on the topic of CSF for ERP projects [15], [16].

For the current analysis, we conducted two separate literature reviews according to the same search procedure and steps. The first one (hereafter referred to as *Review 2017*) was done in mid-2017, and covered the period from mid-2012 until mid-2017. The second review (hereafter referred to as *Review 2023*) was conducted at the end of 2023, covering the period from mid-2017 until the end of 2023. In general, these were database-driven reviews with an additional search in the proceedings of several conferences in the field of information systems (IS). To make our review reproducible, we listed tables with the databases and search terms in the Appendix (see Tables IV and V).

The steps taken in our review procedures are presented in the following paragraphs. An overview is given in Fig. 1 and Fig. 2 with regard to the numbers of papers identified or remaining during/after each step. With each step, the number of papers was reduced according to the assembly of different criteria. In particular, papers were deemed relevant for our analyses that either:

- examined CSFs in empirical studies (qualitative and/or quantitative),
- presented CSFs in literature reviews, or
- derived specific (partly individual) CSFs based on theory.

These criteria were chiefly applied in Step 4 and Step 5 (see below) in the more detailed analysis of the abstracts and paper content.

Step 1: The first step was to define the sources for the literature review. Therefore, several databases and conference proceedings were identified (see Table IV).

Step 2: In this step, we had to define the search terms for the database-driven review. Keywords selected for this search were nearly the same keywords as in our previous studies (see [15], [16]). In those studies, the keywords were mostly derived from those supplied and used by the authors of some of the relevant articles identified in a preliminary literature review. The search terms that we used are listed in Table V.

As not every database has the same search fields and search functions, the search fields listed in Table V differ depending on the database. However, we have made sure that the search fields are at least somewhat similar—insofar as this was possible—depending on the functionalities in the databases.

For not every conference, appropriate search fields or search functionality could be applied. Hence, we decided, if necessary, to review the abstracts and titles of the papers in this step sometimes manually.

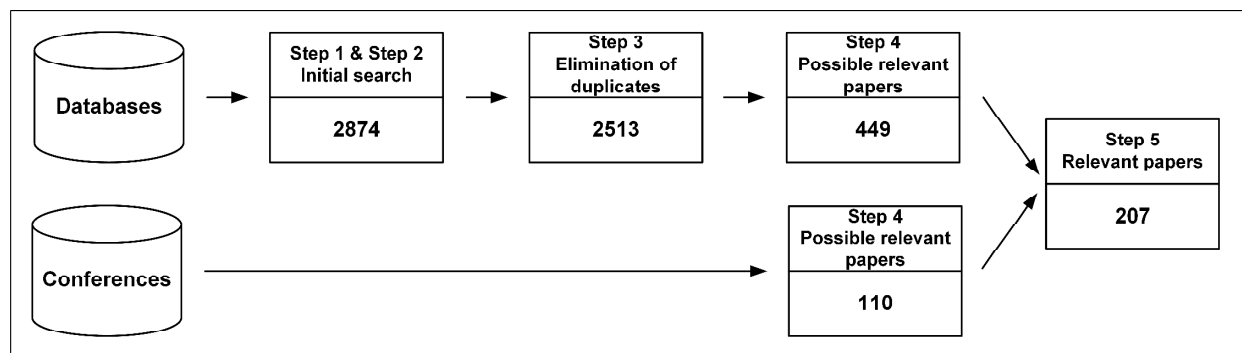


Fig. 1. Progress of Review 2017

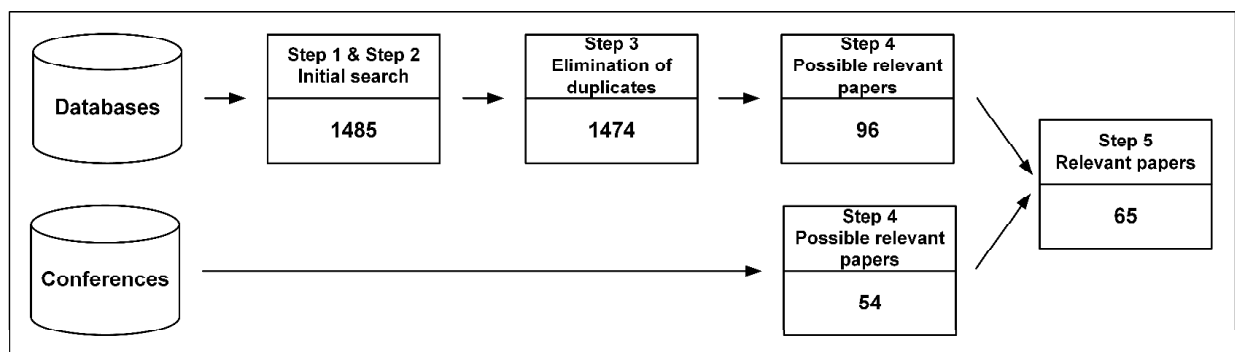


Fig. 2. Progress of Review 2023

Step 3: During Step 3, we performed the initial search according to Steps 1 and 2, and then eliminated duplicates.

- **Review 2017:** The initial search provided 2,874 papers from the databases. After eliminating the duplicates, 2,513 articles remained. From the conference search, 110 papers remained. Altogether, 2,623 papers were identified during the initial search step.
- **Review 2023:** During the initial search step, 1,485 articles were found. After deleting the duplicates, 1,474 papers remained. From the conferences, 54 papers remained. Therefore, altogether 1,528 papers were found during this step.

Step 4: Step 4 included the identification of irrelevant papers. During the initial search, we did not apply any restrictions. The search was not limited to the research field of IS; therefore, papers from other research fields or papers that did not address CSFs of ERP projects (directly or at least indirectly) were also included in the results. Thus, these papers had to be excluded. This was done by reviewing the

abstracts of the papers and, if necessary, by examining the paper content.

- **Review 2017:** Of the identified papers in the previous steps, 449 stemming from the database search and all 110 conference papers remained. In total, this review yielded 559 papers that were potentially relevant to the field of CSFs for ERP system implementations (see Fig. 1).
- **Review 2023:** In this review, 96 papers resulting from the databases and all conference papers remained as potentially relevant. Altogether, 150 had to be read in depth according to Step 5 (see Fig. 2).

Step 5: The fifth and final step consisted of a detailed analysis of the remaining 559 and 150 papers from both reviews and the identification of the CSFs. Therefore, the content of all papers was analyzed in depth to categorize the identified success factors. Emphasis was placed not only on the wording of these factors but on their meaning. Following this step, 207 relevant papers that presented, discussed, or mentioned CSFs remained from Review 2017 and 65 articles remained from Review 2023.

Therefore, both reviews led to a final sum of 272 relevant papers. After identifying these papers and the factors stated within them, we developed a concept matrix to match the factors with the papers for the analysis. For each paper, the CSFs were captured as well as the year, the type of data collection used, and the number and size of the companies from which the CSFs were derived. The results of the analysis of these papers and their respective CSFs are described in the following sections.

III. LITERATURE REVIEWS – RESULTS

A. Critical Success Factors Identified

A CSF for ERP projects has been defined by Finney & Corbett [18] as a reference to any condition or element that was deemed necessary for the ERP project to be successful. Within the analyzed papers, 33 CSFs were identified (see below, ordered alphabetically).

- | | |
|---|---|
| • Available resources (e.g., employees, budget) | • Knowledge management |
| • Balanced project team (cross-functional) | • Monitoring, performance measurement and maintenance |
| • Business process reengineering | • Organizational culture |
| • Change management | • Organizational fit of the ERP |
| • Clear goals and objectives (e.g., business plan, vision, decision strategies) | • Organizational structure |
| • Communication | • Project champion |
| • Company's strategy / strategy fit | • Project leadership / empowered decision makers |
| • Data accuracy (analysis and conversion) | • Project management |
| • Data security and privacy | • Skills, knowledge, and expertise |
| • Environment (e.g., language, culture) | • System and data availability |
| • ERP system acceptance / resistance | • Top management support and involvement |
| • ERP system configuration | • Troubleshooting |
| • ERP system tests | • Use of a steering committee |
| • External consultants | • User training |
| • Interdepartmental cooperation | • Vendor relationship and support |
| • Involvement of end-users and stakeholders | • Vendor tools and implementation methods |
| • IT structure and legacy systems | |

To provide a comprehensive understanding of the different CSFs and their concepts, these are described in this section before presenting a further analysis and discussion of the results. However, only the top five CSFs (see Fig. 3) are described as well as the two factors (*System and data availability* and *Data security and privacy*) that were newly identified in our two current reviews. A detailed description of each individual factor can be requested from the first author or can also be found in Leyh and Sander [16].

Top management support and involvement: Top management support and involvement is one of the most important success factors for ERP projects [19]. A committed leadership at the top management level is the basis for continuous accomplishment of every project [18]. Thus,

innovations, in particular new technologies, are more widely accepted by employees if they are promoted by top management. Before the project starts, top management must identify the peculiarities and challenges of the planned ERP project. Since many decisions that have to be made during the project affect the whole enterprise, they will need the acceptance and the commitment of senior managers, and often they are the only ones who can make such decisions [20]. Commitment of top management is important in allocating necessary resources to make quick and effective decisions, solving conflicts that need enterprise-wide acceptance, and reaching and supporting cooperation from all departments [21].

Organizational fit of the ERP System: The fact that the organizational fit of an ERP system should be examined and considered comprehensively before its implementation appears to be logical. Nevertheless, ERP vendors tend to set up blind confidence in their ERP package even if it is obvious that the organizational fit is weak. Hong and Kim [22] empirically examined to what extent the implementation success of an ERP system depends on the fit between the company and the ERP system and determined that the adaptation and configuration effort negatively correlates with the implementation success. Therefore, it is essential to select an ERP system carefully by considering its specific organizational fit, such as company size or industry sector. Thus, the right ERP system selection is an important factor to ensure the fit between the company and the ERP system.

User training: Often, missing or lacking end-user training is a cause for failure during the implementation of new software. The main goal of end-user training is to provide an effective understanding of the new business processes and applications as well as the new workflows that result from the ERP implementation. Therefore, it is important to set up a suitable plan for the training and education of the employees [21]. Furthermore, during such an extensive project it has to be determined which employee fits best for which position or application of the new software. This depends strongly on the employee's previous knowledge and on who will require additional training courses [23].

Project management: Project management refers to the ongoing management of the implementation plan [18]. The implementation of an ERP system is a unique procedure that requires enterprise-wide project management. Therefore, it involves the planning stages, the allocating of responsibilities, the definition of milestones and critical paths, training and human resource planning, and the determination of measures of success [24], [25]. This enables a better organized approach to decision making, and guarantees that the most suitable company members will make decisions. Furthermore, continuous project management makes it possible to focus on the important aspects of the ERP implementation and ensures timeliness and that schedules are met [24]. Within project management, a comprehensive documentation of the tasks, responsibilities, and goals is indispensable for the success of ERP implementations [26].

Change management: Change management involves early participation of all persons affected by a change process in order to reduce resistance against these modifications. An important component is adequate training, especially of the IT department, as well as early communication of the changes to provide employees with an opportunity to react [24]. Change management strategies are responsible for handling the enterprise-wide cultural and structural changes. Therefore, it is necessary to train and educate the employees in various ways. Change management aims toward preventing

rejection and supporting acceptance. Moreover, its goal is to make employees understand and desire the changes. Early integration of the employees in planning and implementation is important in achieving this understanding. In addition, during the user training sessions, a support team should be available to clarify and answer questions regarding the new processes and function. Furthermore, an additional evaluation with the end-users should be conducted following the "go-live" of the ERP system to uncover problems and avoid discords [27].

Data security and privacy: Data security and privacy play a crucial role in using ERP systems via cloud computing. The migration of sensitive data, as well as the confidentiality of data in the cloud must be guaranteed at all times. Data security also includes the security of services, data centers, and media [28], [29]. Protection against unauthorized access is essential. Data security and privacy are the two main factors involved with cloud computing; however, they also represent concerns on the part of users, which need to be counteracted. Various security techniques and methods can help to protect data in the cloud (see e.g., [30], [31]).

System and data availability: The CSF *System and data availability* describes the capability of a (cloud) ERP system to transmit the right data to the right place at the right time, regardless of location. If there are disruptions and interruptions in accessibility, the implemented system will not fulfill its purpose. Data should be immediately accessible to the end-user via a wide variety of platforms [30], [32]. Especially in times of increasing networking or mobile work, unrestricted availability of services is required. It is, therefore, imperative to agree on continuous availability by means of a service level agreement (SLA) when choosing a provider [33].

B. Analysis of Critical Success Factors and Discussion

As stated above, 272 papers (single- or multiple-case studies, surveys, literature reviews, etc.) were identified that referred to critical success factors of ERP projects. These papers were reviewed again in depth to determine the different concepts of CSFs. Overall, 33 factors (the top five with regard to both of our reviews were described above) were identified. With 33 factors, we used a larger number in our analyses than earlier researchers, because we expected the resulting distribution to be more insightful. If broader definitions of some CSFs are needed later (e.g., grouping the CSFs more coarsely), further aggregation is still possible. Compared to the 31 CSFs from our previous analyses [16], two new factors (also described above) were identified in these reviews: *System and data availability* and *Data security and privacy*. These two new CSFs were identified exclusively in cloud ERP papers.

Table I shows the distribution of all analyzed articles per year. Most of the papers were published between 2012 and 2017. Since 2012 until 2017, approximately more than 30 papers were published each year. Despite similar initial

databases for both reviews, since 2018 only around ten papers or even fewer were published about CSFs each year. Therefore, it could be argued that existing literature reviews should still be updated for the topic area of critical success factors for ERP projects. However, a broader focus should be placed on digitalization projects in general and their CSFs, especially against the backdrop of ever faster developing and changing technologies and technological possibilities.

Fig. 3 shows the results of our reviews in total: the identified CSFs and their overall numbers.

TABLE I. PAPER DISTRIBUTION PER YEAR

Year	Papers	Year	Papers
2023	4	2017	33
2022	7	2016	38
2021	8	2015	44
2020	8	2014	41
2019	12	2013	35
2018	12	2012	30




Fig. 3. CSFs Ordered by Frequency (Number of Articles Analyzed = 272)

Fig. 3 shows that *Top management support and involvement*, *Organizational fit of the ERP system*, *User training*, *Project management* and *Change Management* are the five most frequently named factors, each numbering above 110. The factors *Top management support and involvement* and *Organizational fit of the ERP system* ranked numbers one and two, each of which were referred to in more than 140 papers.

Regarding the form of data collection within the analyzed articles, 24% of the papers covered single- or multiple-case studies, 52% were quantitative surveys (paper-based or online) or qualitative surveys (e.g., as part of interview studies), and 24% of the CSFs were derived on a theoretical basis.

To more specifically analyze and categorize critical success factors, Esteves-Sousa and Pastor-Collado [34] suggested a matrix scheme. They considered the tactical or strategic direction of the CSFs and divided them into organizational and technological factors. Thus, tactical CSFs tended to relate to short-term aspects and goals of the system implementation whereby strategic factors aimed for long-

term impacts of activities with strong connections to the development of the organization in relation to mission, vision and core competencies of the business activity. Considering the technological and organizational character of the CSFs, the specificity and significance of technological factors strongly depended on the ERP systems themselves, whereas organizational factors focused on corporate culture and its environment with its specific processes and structures [34], [35]. Table II gives an overview of the categorization of the top ten CSFs identified in our literature reviews with a focus on their ranking. It is becoming obvious that the majority of important CSFs in ERP projects are organizational factors with a strategic focus. Nevertheless, two of the top five factors are in the area of tactical orientation. Table II also shows that important factors of a technological nature are strategically oriented, which underlines the conscientious selection of the ERP system as essential for the success of the projects. The CSF *IT structure and legacy systems* thereby can take a strategic as well as a tactical perspective. However, technological factors should, therefore, not necessarily be subordinated to organizational considerations.

TABLE II. CATEGORIZATION OF CSFs FOCUSING DIFFERENT PERSPECTIVES (ADAPTED FROM [34], [35])

	Strategic Perspective		Tactical Perspective	
	Critical Success Factor	Rank	Critical Success Factor	Rank
Organizational CSFs	Top management support and involvement	1	User training	3
	Change management	5	Project management	4
	Clear goals and objectives	8		
	Vendor relationship and support	9		
	Balanced project team	10		
Technological CSFs	Organizational fit of the ERP system	2		
	IT structure and legacy systems			6
	ERP system configuration	7		

If we compare the top ten CSFs from Review 2017 and Review 2023 separately (see Table III), it becomes clear that although the ranking of the factors has changed, almost all factors appear in the top ten in both reviews. Only the differentiation of the ranking in Review 2023 based on the frequency of the factors mentioned in the analyzed articles is

not as clear due to the lower number of analyzed articles compared to Review 2017. However, overall, it can be stated that the importance of the factors in general, at least as far as the top ranking is concerned, has not fundamentally changed over time.

TABLE III. COMPARISON OF THE TOP TEN CSFs FROM BOTH REVIEWS (ORDERED BY RANK AND FREQUENCY)

Review 2017			Review 2023		
Rank / Number of instances		CSF	Rank /Number of instances		CSF
1	111	Organizational fit of the ERP system	1	39	Change management
2	109	Top management support and involvement	2	35	Top management support and involvement
3	95	User training	3	33	Project management
4	88	Project management	4	31	User training
5	78	IT structure and legacy systems	4	31	Organizational fit of the ERP system
6	74	ERP system configuration	6	29	IT structure and legacy systems
6	74	Change management	6	29	Business process reengineering
8	72	Clear goals and objectives	8	28	Communication
9	68	Vendor relationship and support	9	26	Vendor relationship and support
9	68	Balanced project team	9	26	ERP system configuration
			9	26	Clear goals and objectives
			9	26	Balanced project team

IV. CONCLUSION AND LIMITATIONS

Regardless of company size, an ERP project is a complex and extensive project that often has a profound impact on company processes. A structured approach to the selection and implementation of ERP systems is, therefore, essential. Taking CSFs into account can be decisive for the success of the project and for a long-term increase in efficiency. Well-thought-out *Project management* in conjunction with other factors such as *Top management support and involvement*, the *Organizational fit of the ERP system*, a *Balanced project team* and comprehensive *User training* make a significant contribution to ensuring that the ERP system is optimally adapted to the specific requirements of the company. Effective *Change management* helps to overcome resistance within the company and ensure smooth integration.

By paying attention to CSFs, companies can not only minimize project risks, but also ensure that the ERP system brings sustainable benefits, such as improved processes, increased transparency, and enhanced competitiveness.

Well-planned ERP projects also enable optimized use of existing resources and seamless integration into the company's existing ES landscape. The implementation of a

suitable ERP system as the basis of the ES landscape ensures that the company is also able to meet future challenges. The successful implementation and use of an ERP system forms an essential building block for the integration of new technologies and lays the foundation for the company's future digitalization initiatives. Both organizational and technological success factors play a decisive role, and these should be carefully coordinated in order to contribute to the success of the ERP project.

Future research activities in this area can build on the insights gained from our studies. Within our own research project, we plan to investigate individual CSFs (especially the most important ones) in more detail to derive and update recommendations for action for the best use of the CSFs in ERP projects. We plan to apply a qualitative approach by conducting several in-depth interviews with different enterprises. Furthermore, it could be an interesting and informative viewpoint to conduct a study that considers the specifics of individual industry sectors and a more specific focus on company sizes. In this way, we hope to further detail the importance of CSFs for ERP projects. Another starting point for future research could be to analyze CSFs with reference to the different types of ERP systems (on premise

vs. cloud) and the different types of ERP projects (implementation vs. adjustments) to investigate differences in the importance and the definitions of the CSFs. Furthermore, a focus of future study should regard the difficulties of implementing individual CSFs in companies and how these obstacles can be minimized.

Regarding our literature review procedures, there are limitations that must be mentioned. We are aware that we cannot be certain that we have identified all relevant papers published in journals and conferences, since we made a selection of specific databases and conferences. Therefore, journals not included in our databases and proceedings of other conferences might also comprise relevant articles. Another limitation is the coding of the CSFs. We tried to reduce subjectivity by formulating coding rules (based on the approach of our previous literature studies) and by discussing the coding of the CSFs among three independent researchers. Hence, other researchers may code the CSFs in another way.

REFERENCES

- [1] R. Hentschel and C. Leyh, "Cloud Computing: Status quo, aktuelle Entwicklungen und Herausforderungen," in *Cloud Computing*, S. Reinheimer, Ed., Wiesbaden: Springer, 2018, pp. 3–20. doi: 10.1007/978-3-658-20967-4_1.
- [2] L. Markfort et al., "Patterns of business model innovation for advancing IoT platforms," *Journal of Service Management*, vol. 33, no. 1, pp. 70–96, 2022, doi: 10.1108/JOSM-11-2020-0429.
- [3] H. Gebauer et al., "How to convert digital offerings into revenue enhancement – Conceptualizing business model dynamics through explorative case studies," *Industrial Marketing Management*, vol. 91, pp. 429–441, 2020, doi: 10.1016/j.indmarman.2020.10.006.
- [4] M. Jovanović, A.-L. Mesquida, A. Mas, and B. Lalić, "Towards the Development of a Sequential Framework for Agile Adoption," in *Software Process Improvement and Capability Determination*, vol. 770, A. Mas, A. Mesquida, R. V. O'Connor, T. Rout, and A. Dorling, Eds., Cham: Springer International Publishing, 2017, pp. 30–42. doi: 10.1007/978-3-319-67383-7_3.
- [5] C. Leyh, W. Kusturica, S. Neuschl, and C. Laroque, "Zukunfts- und Wertschöpfungslabor DataLab WestSax – Ein regionaler Katalysator für datenbasierte Wertschöpfungsprozesse," *Industrie 4.0 Management*, vol. 2022, no. 6, pp. 37–41, 2022, doi: 10.30844/IM_22-6_37-41.
- [6] C. Leyh and K. Bley, "Digitalisierung: Chance oder Risiko für den deutschen Mittelstand? – Eine Studie ausgewählter Unternehmen," *HMD Praxis der Wirtschaftsinformatik*, vol. 53, no. 1, pp. 29–41, 2016, doi: 10.1365/s40702-015-0197-2.
- [7] O. Gassmann and P. Sutter, Eds., *Digitale Transformation im Unternehmen gestalten: Geschäftsmodelle, Erfolgsfaktoren, Handlungsanweisungen, Fallstudien*. München: Hanser, 2016.
- [8] S. Mathrani, A. Mathrani, and D. Viehland, "Using enterprise systems to realize digital business strategies," *Journal of Enterprise Information Management*, vol. 26, no. 4, pp. 363–386, 2013, doi: 10.1108/JEIM-01-2012-0003.
- [9] C. Leyh and T. Schäffer, "Digitale Kompetenzen als notwendige Voraussetzung der Digitalen Transformation," *HMD Praxis der Wirtschaftsinformatik*, vol. 61, no. 1, pp. 12–26, 2024, doi: 10.1365/s40702-024-01044-9.
- [10] M. Pagani, "Digital Business Strategy and Value Creation: Framing the Dynamic Cycle of Control Points," *MIS Quarterly*, vol. 37, no. 2, pp. 617–632, 2013, doi: 10.25300/MISQ/2013/37.2.13.
- [11] C. Leyh, K. Bley, and M. Ott, "Chancen und Risiken der Digitalisierung: Befragungen ausgewählter KMU," in *Arbeit 4.0 – Digitalisierung, IT und Arbeit*, J. Hofmann, Ed., Wiesbaden: Springer, 2018, pp. 29–51. doi: 10.1007/978-3-658-21359-6_3.
- [12] C. Leyh, T. Schäffer, K. Bley, and S. Forstehäusler, "Assessing the IT and Software Landscapes of Industry 4.0-Enterprises: The Maturity Model SIMMI 4.0," in *Information Technology for Management: New Ideas and Real Solutions*, Lecture Notes in Business Information Processing (LNBIP), vol. 277, E. Ziemba, Ed., Cham: Springer, 2017, pp. 103–119. doi: 10.1007/978-3-319-53076-5_6.
- [13] J. Schlick, P. Stephan, M. Loskyll, and D. Lappe, "Industrie 4.0 in der praktischen Anwendung," in *Industrie 4.0 in Produktion, Automatisierung und Logistik*, T. Bauernhansl, M. Ten Hompel, and B. Vogel-Heuser, Eds., Wiesbaden: Springer, 2014, pp. 57–84. doi: 10.1007/978-3-658-04682-8_3.
- [14] R. Winter, B. Bender, and S. Aier, "Enterprise-Level IS Research – Need, Conceptualization, Exemplary Knowledge Contributions and Future Opportunities," in *Proceedings of the 57th Hawaii International Conference on System Sciences (HICSS 2024)*, Honolulu, Hawaii, USA, 2024.
- [15] C. Leyh, "Critical Success Factors for ERP System Implementation Projects: A Literature Review," in *Advances in Enterprise Information Systems II*, C. Möller and S. Chaudhry, Eds., Leiden, The Netherlands: CRC Press/Balkema, 2012, pp. 45–56.
- [16] C. Leyh and P. Sander, "Critical Success Factors for ERP System Implementation Projects: An Update of Literature Reviews," in *Enterprise Systems. Strategic, Organizational, and Technological Dimensions*, Lecture Notes in Business Information Processing (LNBIP), vol. 198, D. Sedera, N. Gronau, and M. Sumner, Eds., Cham: Springer, 2015, pp. 45–67. doi: 10.1007/978-3-319-17587-4_3.
- [17] J. Webster and R. T. Watson, "Analyzing the Past to Prepare for the Future: Writing a Literature Review," *MIS Quarterly*, vol. 26, no. 2, pp. 13–23, 2002.
- [18] S. Finney and M. Corbett, "ERP implementation: a compilation and analysis of critical success factors," *Business Process Management Journal*, vol. 13, no. 3, pp. 329–347, 2007, doi: 10.1108/14637150710752272.
- [19] P. Achanga, E. Shehab, R. Roy, and G. Nelder, "Critical success factors for lean implementation within SMEs," *Journal of Manufacturing Technology Management*, vol. 17, no. 4, pp. 460–471, 2006, doi: 10.1108/17410380610662889.
- [20] J. Becker, O. Vering, A. Winkelmann, and M. Bartsch, Eds., *Softwareauswahl und -einführung in Industrie und Handel: Vorgehen bei und Erfahrungen mit ERP- und Warenwirtschaftssystemen*. Berlin Heidelberg: Springer, 2007.
- [21] M. Al-Mashari, A. Al-Mudimigh, and M. Zairi, "Enterprise resource planning: A taxonomy of critical factors," *European Journal of Operational Research*, vol. 146, no. 2, pp. 352–364, 2003, doi: 10.1016/S0377-2217(02)00554-4.
- [22] K.-K. Hong and Y.-G. Kim, "The critical success factors for ERP implementation: an organizational fit perspective," *Information & Management*, vol. 40, no. 1, pp. 25–40, 2002, doi: 10.1016/S0378-7206(01)00134-3.
- [23] I. Teich, W. Kolbenschlag, and W. Reiners, *Der richtige Weg zur Softwareauswahl*. Berlin, Heidelberg: Springer, 2008. doi: 10.1007/978-3-540-71262-6.
- [24] M. Al-Mashari and A. Al-Mudimigh, "ERP implementation: lessons from a case study," *Information Technology & People*, vol. 16, no. 1, pp. 21–33, 2003, doi: 10.1108/09593840310463005.
- [25] F. Fui-Hoon Nah, J. Lee-Shang Lau, and J. Kuang, "Critical factors for successful implementation of enterprise systems," *Business Process Management Journal*, vol. 7, no. 3, pp. 285–296, 2001, doi: 10.1108/14637150110392782.
- [26] B. Snider, G. J. C. Da Silveira, and J. Balakrishnan, "ERP implementation at SMEs: analysis of five Canadian cases," *International Journal of Operations & Production Management*, vol. 29, no. 1, pp. 4–29, 2009, doi: 10.1108/01443570910925343.
- [27] T. C. Loh and S. C. L. Koh, "Critical elements for a successful enterprise resource planning implementation in small-and medium-sized enterprises," *International Journal of Production Research*, vol. 42, no. 17, pp. 3433–3455, 2004, doi: 10.1080/00207540410001671679.
- [28] A. Alharthi, M. O. Alassafi, R. J. Walters, and G. B. Wills, "An exploratory study for investigating the critical success factors for cloud migration in the Saudi Arabian higher education context," *Telematics and Informatics*, vol. 34, no. 2, pp. 664–678, 2017, doi: 10.1016/j.tele.2016.10.008.
- [29] S. R. Tehrani and F. Shirazi, "Factors Influencing the Adoption of Cloud Computing by Small and Medium Size Enterprises (SMEs)," in *Human Interface and the Management of Information. Information and*

- Knowledge in Applications and Services*, Lecture Notes in Computer Science (LNCS), vol. 8522, S. Yamamoto, Ed., Cham: Springer, 2014, pp. 631–642. doi: 10.1007/978-3-319-07863-2_60.
- [30] T. N. Mahara, “Indian SMEs Perspective for election of ERP in Cloud,” *Journal of International Technology and Information Management*, vol. 22, Article 5, 2013, doi: 10.58729/1941-6679.1004.
- [31] Y. Sun, J. Zhang, Y. Xiong, and G. Zhu, “Data Security and Privacy in Cloud Computing,” *International Journal of Distributed Sensor Networks*, vol. 10, no. 7, 2014, doi: 10.1155/2014/190903.
- [32] M. H.L., A. O. Mathew, and L. L. R. Rodrigues, “Prioritizing the factors affecting cloud ERP adoption – an analytic hierarchy process approach,” *International Journal of Emerging Markets*, vol. 13, no. 6, pp. 1559–1577, 2018, doi: 10.1108/IJoEM-10-2017-0404.
- [33] M. Lechesa, L. Seymour, and J. Schuler, “ERP Software as Service (SaaS): Factors Affecting Adoption in South Africa,” in *Re-conceptualizing Enterprise Information Systems*, Lecture Notes in Business Information Processing (LNBIP), vol. 105, C. Möller and S. Chaudhry, Eds., Berlin, Heidelberg: Springer, 2012, pp. 152–167. doi: 10.1007/978-3-642-28827-2_11.
- [34] J. Esteves-Sousa and J. Pastor-Collado, “Towards the Unification of Critical Success Factors for ERP Implementations,” in *Proceedings of the 10th Annual Business Information Technology Conference (BIT 2000)*, Manchester, UK, 2000.
- [35] U. Remus, “Critical success factors for implementing enterprise portals: A comparison with ERP implementations,” *Business Process Management Journal*, vol. 13, no. 4, pp. 538–552, 2007, doi: 10.1108/14637150710763568.

APPENDIX

TABLE IV. SOURCES FOR THE LITERATURE REVIEW

Databases	Conferences
<ul style="list-style-type: none"> • AIS eLibrary (only Review 2023) • Academic Search Complete and Business Source Complete (both only Review 2017) • Business Source Premier (only Review 2023) • Emerald Insight • Science Direct • SpringerLink 	<ul style="list-style-type: none"> • AMCIS • Bled eConference • ECIS • HICCS • ICIS • PACIS • Wirtschaftsinformatik (WI)

TABLE V. SEARCH FIELDS AND SEARCH TERMS

Database + Search fields	Search terms / Keywords
AIS eLibrary: “Title” or “Abstract” or “Subject”	<ul style="list-style-type: none"> • ERP + success* • ERP + failure
Business Source Premier & Academic Search Complete & Business Source Complete: “TI Title” or “AB Abstract”	<ul style="list-style-type: none"> • ERP + crit* • ERP + CSF • ERP + CFF
Emerald Insight: “Abstract”	<ul style="list-style-type: none"> • ERP + fact* • "Enterprise system*" + success*
Science Direct: “Title, Abstract, Keywords”	<ul style="list-style-type: none"> • "Enterprise system*" + failure • "Enterprise system*" + crit*
SpringerLink: “where the title contains”	<ul style="list-style-type: none"> • "Enterprise system*" + CSF • "Enterprise system*" + CFF • "Enterprise system*" + fact*

d'Alembert Convolution for Enhanced Spatio-Temporal Analysis of Forest Ecosystems

Rytis Maskeliūnas
CoE Forest 4.0
Vytautas Magnus University
Kaunas, Lithuania
rytis.maskeliunas@vdu.lt

Robertas Damaševičius
CoE Forest 4.0
Kaunas University of Technology
Kaunas, Lithuania
robertas.damasevicius@ktu.lt

Abstract—This paper presents a novel approach to enhance the spatio-temporal analysis of forest ecosystems using the d'Alembert convolution method, which, integrating elements from wave equation theory and convolutional neural networks, enables the comprehensive analysis of remote sensing images by capturing both spatial and temporal variations. This methodology not only improves feature extraction, but also helps address the challenges associated with traditional image processing techniques, which often overlook the temporal dynamics of forests. The results show significant improvements in the analysis of forest ecosystems. Specifically, the higher performance metrics compared to existing methods, including higher accuracy in classifying various forest types and more effective monitoring of changes over time.

Index Terms—Convolutional Neural Networks, Spatio-Temporal Analysis, Remote Sensing, d'Alembert Operator.

I. INTRODUCTION

THE IMPORTANCE of forests in global environmental health, biodiversity conservation, and economic resources is undeniable. Forests play a crucial role in carbon sequestration, climate regulation, and providing habitats for a variety of species [1]. However, they are constantly under threat from various factors, including climate change, deforestation, pests, and diseases [2]. Traditional forest management methods, often relying on periodic and manual surveys, are inadequate in the face of these rapidly evolving challenges [3]. The need for more efficient, accurate, and real-time monitoring and management methods is more pressing than ever [4].

Forestry management is entering a new era of technological innovation, marked by the integration of advanced computational methods and environmental science. The advent of smart forestry, using data-driven approaches, has opened new pathways for sustainable forest management and environmental conservation [5]. Advancements in remote sensing technologies, such as satellite imagery and aerial photography, have propelled the field of forestry into the digital age [6]. Hyperspectral imaging, in particular, has become a valuable tool for monitoring vegetation health, biomass estimation, and detecting changes within forest ecosystems [7]. However, the sheer volume and complexity of the generated data pose significant challenges in terms of processing and analysis [8]. Conventional image processing techniques, while beneficial,

often fail in extracting the full spectrum of information hidden within multidimensional spatial and temporal data sets [9]. However, the complexity and dynamic nature of forest ecosystems pose great challenges in terms of data collection, analysis, and interpretation [10], one of the challenges still present in smart forest, as outlined in [11].

Convolutional neural networks (CNNs) have revolutionized the field of image analysis [12]. However, their application in forestry has been somewhat limited, focusing mainly on spatial data without fully exploiting the temporal dimension [13]. By extending the convolution operation to incorporate differential operators that account for both spatial and temporal changes, similar to the components of the d'Alembert operator, we propose a novel method that not only enhances feature extraction from remote sensing images, but also captures the dynamic changes occurring within forest ecosystems over time by adapting the convolution operation to include differential operators that account for spatial and temporal changes, this method allows for a more nuanced extraction of features from remote sensing images, surpassing the capabilities of traditional convolutional methods. The approach allows for an increased level of detail in analyzing both spatial and temporal variations in forest ecosystems, contributing to a deeper understanding of forest dynamics.

II. RELATED WORKS

Existing research has demonstrated a robust exploration of remote sensing applications, taking advantage of advanced machine learning and deep learning techniques to address a spectrum of challenges in land use and cover change (LULC), environmental monitoring and resource management. Several researchers utilize machine learning models to analyze and interpret LULC changes and their impacts on ecosystems. For example, Saha et al. (2024) employ geospatial techniques and machine learning to assess the degradation of the Deepor wetland in India, highlighting high losses due to urbanization and agricultural expansion [14]. Similarly, Thien et al. (2023) examined the spatiotemporal dynamics of LULC in Vietnam's Red River delta, attributing changes predominantly to urban development [15]. In a broader scope, Masolele et al. (2021) deploy spatial and temporal deep learning methods to classify land use after tropical deforestation, underscoring the supe-

rior performance of spatio-temporal models over conventional approaches [16]. Mareto et al. (2021) further this discourse by mapping deforestation in the Amazon, demonstrating how spatio-temporal deep learning improves monitoring accuracy [17]. Others focused on the development of sophisticated machine learning algorithms to refine remote sensing data retrieval and analysis. Fonseca et al. (2023) innovate in multi-temporal SAR image analysis through wavelet spatio-temporal change detection, achieving high accuracy with reduced computational demands [18]. On a similar note, Dimiyati et al. (2023) and Jing et al. (2023) introduce methods to monitor mangrove changes and combine remote sensing images, respectively, showcasing the potential of these advanced techniques in managing complex environmental datasets [19], [20].

III. THEORETICAL FOUNDATIONS

A. Convolutional Neural Networks in Image Processing

Convolutional Neural Networks (CNNs) have revolutionized the field of image processing and computer vision. They are a specialized kind of neural network designed for processing data with a grid-like topology, such as images. A CNN learns to recognize patterns and features in images through the process of convolution, pooling, and fully connected layers.

The convolution of a function f with a kernel g is defined as:

$$(F * G)(x, y) = \sum_{i=-a}^a \sum_{j=-b}^b F(i, j) \cdot G(x - i, y - j) \quad (1)$$

where F is the image, G is the kernel, and x, y are spatial coordinates in the image, and a and b represent the half-width and half-height of the kernel G , respectively. The kernel G slides over the image F , computing the sum of element-wise products at each position.

CNN architecture has several types of layers. Lower layers capture basic features such as edges and textures, while deeper layers identify complex patterns specific to the training data:

- **Convolutional Layer** performs the convolution operation. It applies a set of learnable filters (kernels) to the input image. Each filter extracts different features from the input.

$$C_{out} = \text{ReLU}(C_{in} * K + b) \quad (2)$$

where C_{in} is the input, K is the convolutional kernel, b is a bias term, and ReLU is the activation function, typically a Rectified Linear Unit.

- **Pooling Layer** reduces the spatial size (height and width) of the input volume, making the network computation more efficient. It also helps to make the network invariant to small translations of the input.
- **Fully Connected Layer** has connections to all activations in the previous layer. These layers are typically used at the end of the network to perform classification based on the features extracted by the convolutional and pooling layers.

B. The d'Alembert Operator in Wave Equations

The d'Alembert operator is a second-order differential operator that is also widely used in the fields of physics and engineering, particularly in the study of wave propagation and vibrations. The d'Alembert operator is defined in the context of a four-dimensional space-time continuum, combining time and space derivatives. In a three-dimensional space with time, the operator is represented as:

$$\square = \frac{\partial^2}{\partial t^2} - \nabla^2 \quad (3)$$

where $\frac{\partial^2}{\partial t^2}$ is the second derivative with respect to time, and ∇^2 is the Laplacian operator, which is a scalar differential operator defined as the divergence of the gradient of a function, representing the sum of second spatial derivatives:

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \quad (4)$$

The wave equation is particularly suitable for analyzing spatiotemporal data in remote sensing images because it inherently captures the propagation of information over both space and time, which aligns with the dynamic nature of environmental phenomena. It combines the second-order temporal derivative with the Laplacian operator, effectively linking temporal changes to spatial variations. This dual capability allows the wave equation to model how disturbances, such as changes in vegetation or land cover, propagate through time and space, providing a comprehensive framework for tracking these dynamics. By incorporating the d'Alembert operator into convolutional analysis, the network can utilize these properties to enhance feature extraction, capturing both the immediate spatial details and their evolution over time. This results in a more robust analysis of remote sensing data, as it allows for the detection and interpretation of complex temporal patterns and spatial structures within the forest ecosystem, essential for accurate monitoring and assessment. The wave equation for a scalar field $\psi(x, y, z, t)$ in a three-dimensional space can be written as:

$$\square\psi = \frac{\partial^2\psi}{\partial t^2} - \nabla^2\psi = 0 \quad (5)$$

This equation describes how a wave propagates in space and time. The term $\frac{\partial^2\psi}{\partial t^2}$ represents the acceleration of the wave, while $\nabla^2\psi$ accounts for the spatial spread of the wave.

C. d'Alembert Operator for Convolutional Analysis

CNNs handle spatial data by applying convolutional filters to extract features such as edges, textures, and patterns from static images. These filters are applied across the spatial dimensions (height and width) of the image, capturing local spatial hierarchies and invariances. Temporal processing, on the other hand, works by capturing the changes and dynamics over time, which is critical for understanding phenomena like forest growth or seasonal variations in remote sensing data. To achieve this, the network incorporates layers or mechanisms that can capture temporal dependencies - the d'Alembert operator, which, traditionally used to describe wave propagation,

is adapted to account for both spatial and temporal changes by modifying the convolution operation to include differential components. By using a kernel enhanced with the d'Alembert operator, the network simultaneously processes the spatial features (through the traditional convolution) and the temporal features (through the operator's temporal derivative component). Therefore, this integration allows the network to learn more comprehensive feature representations that encapsulate both static and dynamic aspects of the data, providing a robust and versatile framework for tasks requiring spatio-temporal analysis. The balance is achieved by including both standard convolutional layers for spatial processing and modified layers to handle temporal dynamics effectively.

To integrate the d'Alembert operator into the convolution process, we first redefine the convolution operation to include differential components. Consider a remote sensing image sequence represented as $I(x, y, t)$, where (x, y) are spatial coordinates, and t is the time dimension. The adapted convolution operation, incorporating the d'Alembert operator, can be mathematically represented as:

$$C_{d'Alembert}^l(x, y, t) = (K * (\square I))(x, y, t) \quad (6)$$

where K is the convolution kernel, $*$ denotes the convolution operation, and \square is the d'Alembert operator. The d'Alembert operator applied to the image sequence I is defined as:

$$\square I(x, y, t) = \frac{\partial^2 I}{\partial t^2}(x, y, t) - \nabla^2 I(x, y, t) \quad (7)$$

The Laplacian component $\nabla^2 I(x, y, t)$ of the d'Alembert operator enhances the extraction of spatial features by emphasizing areas with high spatial frequency, such as edges and textures in the image.

The temporal derivative component $\frac{\partial^2 I}{\partial t^2}(x, y, t)$ captures changes in the image sequence over time, highlighting the dynamic changes in the forest environment.

The integration of the d'Alembert convolution into a CNN framework necessitates a modification of the standard convolutional layers. This modification involves applying a kernel that is enhanced with the d'Alembert operator, enabling the network to capture both spatial and temporal variations more effectively. The d'Alembert-enhanced kernel is designed to incorporate both the spatial features, captured by the traditional convolution kernel, and the temporal features, introduced by the d'Alembert operator. Consider a standard convolution kernel K and its adaptation with the d'Alembert operator.

$$K_{d'Alembert}(x, y, t) = K(x, y) + \lambda \cdot (\square I)(x, y, t) \quad (8)$$

Here, $K(x, y)$ is the standard convolution kernel, $\square I$ represents the application of the d'Alembert operator on the image sequence I , and λ is a weighting factor that balances the spatial and temporal components.

The convolution operation in a CNN is modified to use this d'Alembert-enhanced kernel. The modified convolution operation for an input image sequence I at layer l is:

$$C_{d'Alembert}^l(x, y, t) = (K_{d'Alembert}^l * I^l)(x, y, t) \quad (9)$$

where $*$ denotes the convolution operation, and l indicates the layer in the CNN.

The d'Alembert convolution allows the CNN to extract features that encapsulate both spatial variations (such as edges, textures) and temporal changes (such as growth patterns, environmental dynamics). This dual capability is required for the analysis of remote sensing images of forests, where both spatial and temporal indicators are required to understand forest health and dynamics.

$$F_{d'Alembert}^l = \text{Activation}(C_{d'Alembert}^l(x, y, t)) \quad (10)$$

Here, $F_{d'Alembert}^l$ represents the feature maps obtained after applying the d'Alembert convolution at layer l , and Activation denotes the activation function used in the CNN (e.g., ReLU).

D. Architecture of d'Alembert Network

The architecture of the d'Alembert network is presented in Table I and is discussed in detail below.

TABLE I
ARCHITECTURE OF THE D'ALEMBERT NETWORK

Layer Type	Output Size	Kernel Size	Other Parameters
Input Layer	480x480x3	-	-
Convolutional Layer	470x470x32	11x11	Stride=1, Padding=Valid
Activation Layer	470x470x32	-	ReLU
Pooling Layer	235x235x32	2x2	Stride=2, Type=Max
d'Alembert Conv Layer	225x225x64	11x11	Stride=1, Padding=Valid, $\lambda = 0.1$
Activation Layer	225x225x64	-	ReLU
Pooling Layer	112x112x64	2x2	Stride=2, Type=Max
Fully Connected Layer	1024	-	-
Activation Layer	1024	-	ReLU
Fully Connected Layer	512	-	-
Activation Layer	512	-	ReLU
Output Layer	7	-	Softmax

The basis of the d'Alembert network is of convolutional layers for extracting spatial features from imagery. Each convolutional layer applies a set of learnable filters to the input image, detecting features such as edges, textures, and shapes. These features are needed for the structural components of the forest, such as the canopy density and tree boundaries. The convolution operation combines image data with a kernel (filter) through a dot product that aggregates local pixel values to produce a feature map, highlighting areas of interest in the image. Following each convolutional layer, an activation function is used (the Rectified Linear Unit (ReLU)). ReLU introduces non-linearity into the model, allowing it to learn more complex patterns. It works by replacing all negative pixel values in the feature map with zero, maintaining only positive values that correspond to detected features, and it helps to overcome the problem of vanishing gradients, ensuring that the network continues to learn effectively throughout its depth. The inclusion of pooling layers (max pooling), reduces the spatial size of the representation, making the computation more manageable, and the network more robust to variations

in the image. By downsampling the feature maps, the pooling layers help reduce the amount of data that needs to be processed while preserving the most essential information, such as the dominant features within a local patch of the image. Which is particularly useful in forest imagery, where specific features, such as tree clusters and clearings, need to be emphasized over large, uniform areas.

The main novelty in our d'Alembert network is the adaptation of the d'Alembert operator into the convolution process. This operator is applied here to account for both spatial and temporal changes in forest imagery by modifying the convolution operation to include differential operators (Laplacian for spatial and second-order time derivative for temporal features). In this way, the network can capture dynamic changes in the forest, such as growth, deforestation, or seasonal variations. Toward the end of the network, fully connected layers are used to interpret the features extracted and learned by the convolutional and pooling layers. These layers consolidate the learned features into a format suitable for classification or regression tasks, such as identifying different types of forests or assessing forest health. Each neuron in these layers connects to all activations in the previous layer, allowing the network to learn non-linear combinations of the high-level features. The final layer of the network (softmax layer) outputs a probability distribution over the target classes. For forest imagery analysis, these classes include different types of land cover, such as dense forest, degraded forest, water bodies, etc. The softmax function converts the logits from the fully connected layer into probabilities by exponentiating and normalizing each output, providing a clear, interpretable classification result.

The performance sensitivity to hyperparameters such as learning rate, batch size, and the number of convolutional layers is unfortunately quite significant in our approach. The learning rate strongly affects the model's ability to effectively integrate both spatial and temporal features; a high learning rate lead to suboptimal convergence in the complex landscape of spatio-temporal data, while a low rate cause excessively slow training, missing critical temporal patterns. Batch size influences the network's capacity to generalize from dynamic forest data; larger batch sizes provide more stable gradient estimates, improving convergence and capturing broader temporal changes, but at the cost of higher memory usage. Conversely, smaller batches enhance generalization but introduce noisy gradients, potentially destabilizing training. The number of convolutional layers directly impacts the depth of spatial feature extraction; insufficient layers fail to capture the intricate textures and edges within forest images, while too many layers could overfit the spatial details and neglect temporal dynamics. Hyperparameter tuning is therefore required to accurately detect and analyze both spatial and temporal variations essential for monitoring forest ecosystems.

IV. CASE STUDY AND EXPERIMENTAL RESULTS

A. Datasets

The DeepGlobe Land Cover 2018 dataset [21] is a collection of high-resolution satellite images used for land cover

classification challenges, focusing on categorizing land cover into multiple classes. It encompasses geographical landscapes from different parts of the world, offering a robust platform to advance land cover analysis technologies. The dataset consists of 1146 images with 3042 labeled objects belonging to 7 different classes including agriculture_land, urban_land, rangeland, water, barren_land, forest_land, and unknown (see sample images in Figure 1). In this study, we used a subset of the DeepGlobe dataset, DeepGlobe-Forest, which includes only 191 images labeled as forest_land.

The LoveDA [22], [23] dataset is a remote sensing dataset adapted for the study of natural landscapes and their dynamic changes. It consists of multispectral imagery collected from various satellite platforms. The images in the dataset have high spatial resolution, which aids in detailed analysis and facilitates accurate monitoring of small-scale changes. It includes a mix of urban and rural landscapes for a diverse range of scenes. The dataset consists of 5987 images with 20658 labeled objects belonging to seven different classes, including background, road, building, forest, water, agriculture, and barren (see sample images in Figure 2). We used a subset of the LoveDA dataset, called LoveDA-Forest, which has 3043 images labeled as forest. Both the DeepGlobe-Forest and LoveDA-Forest datasets have been used previously in [24]. Preprocessing involved resizing images to a consistent dimension, normalizing pixel values to a standard scale, and augmenting the data with transformations such as rotations and flips to enhance model generalization. Additionally, temporal alignment of sequential images was ensured for the LoveDA dataset to capture temporal dynamics accurately.

TABLE II
COMPARISON OF DEEPGLOBE-FOREST AND LOVEDA-FOREST DATASETS

Characteristic	DeepGlobe-Forest	Loveda-Forest
Sensor	DigitalGlobe	Sentinel-2
Image Size	2448x 2448	1024 x 1024
Spectral Range (μm)	0.4 - 2.3	0.45 - 2.4
Number of Bands	11	13
Spatial Resolution (m)	30	10
Number of Classes	5	7

B. Experimental setting

For this study, the analysis was performed using a custom implementation developed in TensorFlow. The model training used the DeepGlobe-Forest and Loveda-Forest datasets in 200,000 epochs, starting with an initial learning rate of 0.00005. This rate was progressively reduced starting from the 2000th epoch using a linear decay strategy. The input images were cropped to a uniform size of 480×480 for consistency in both datasets. The computational resources included a single NVIDIA RTX 2060 graphics card, an AMD Ryzen 9 5950X CPU and 32 GB of RAM.

C. Performance evaluation

We used performance metrics such as overall accuracy (OA), average accuracy (AA) per class, and the Kappa coefficient (Kappa). The OA metric reflects the proportion of



Fig. 1. Sample images from DeepGlobe dataset: water, barren_land, forest_land, urban_land, and agriculture_land



Fig. 2. Sample images from LoveDA dataset: forest, barren, water, agriculture, background, road, and building

correctly classified images in the test dataset relative to the total sample count. The AA metric represents the average accuracy in each image class, while the Kappa metric provides a measure of accuracy adjusted for the probability of random chance. In addition, precision, recall, and F-1 score were also utilized as performance indicators. The results are summarized in table III.

TABLE III
PERFORMANCE METRICS FOR DEEPGLOBE-FOREST AND
LOVEDA-FOREST DATASETS

OA (%)	AA (%)	Kappa	Precision(%)	Recall(%)	F1 Score(%)
DeepGlobe-Forest					
91.19	83.89	0.88	92.13	90.24	91.17
LoveDA-Forest					
86.89	74.34	0.84	87.31	85.45	88.37

For the DeepGlobe-Forest dataset, achieving an OA of 91.19% indicates high accuracy in classifying forest cover types, supported by an AA of 83.89% per class and a robust Kappa coefficient of 0.88, demonstrating strong statistical result reliability. Precision scores around 92.13% and balanced recall rates of 90.24% with a high F1 score of 91.17%, show that the model is able to accurately identify and differentiate forest categories. Similarly, the LoveDA-Forest dataset shows an OA of 86.89%, with precise classification reflected in precision and recall scores of 87.31% and 85.45%, respectively, and an F1 score of 88.37%.

D. Results of the segmentation performance

The segmentation performance of the models was evaluated using two metrics: Intersection over Union (IoU) and Accuracy (Acc):

$$\text{IoU} = \frac{p_{ii}}{\sum_{j=0}^k p_{ij} + \sum_{j=0}^k p_{ji} - p_{ii}},$$

where p_{ij} denotes the prediction of the category i into category j , and $k + 1$ is the total number of categories. The mean Intersection over Union (mIoU) is calculated by:

$$mIoU = \frac{1}{k+1} \sum_{i=0}^k \frac{p_{ii}}{\sum_{j=0}^k p_{ij} + \sum_{j=0}^k p_{ji} - p_{ii}},$$

The formulas for the accuracy and mean accuracy (mAcc) are given by:

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + FN + TN},$$

$$mAcc = \frac{1}{k+1} \sum_{i=0}^k \frac{TP_i + TN_i}{TP_i + FP_i + FN_i + TN_i},$$

where TP represents the true positives, TN the true negatives, FP the false positives, FN the false negatives, and $k + 1$ indicates the total number of categories.

For comparison we include the values of deeplabv3+ [25], pidnet [26], pspnet [27], knet [28], segformer [29], mask2former [30] and segnext [31] as was determined in the research of Wang et al. [24] (we have not replicated these methods in our work).

TABLE IV
COMPARISON OF MODEL PERFORMANCES ON FOREST AND BACKGROUND CLASSES FROM DEEPGLOBE-FOREST DATASET [24].

Model	IoU (%)		mIoU (%)		Accuracy (%)	
	Forest	Back-ground	Forest	Back-ground	Forest	Back-ground
Deeplabv3+	77.69	79.67	78.68	87.42	88.70	88.06
Segformer	79.98	81.71	80.85	89.02	89.80	89.41
Pidnet-s	78.78	80.96	79.87	87.35	90.21	88.78
Mask2former	80.52	81.61	81.06	91.09	88.17	89.63
Pspnet	79.86	80.79	80.33	91.17	87.22	89.20
Segnext	80.60	81.84	81.22	90.69	88.71	89.70
Knet-s3-r50	80.23	81.22	80.73	91.24	87.63	89.44
SegForest	82.80	83.99	83.39	91.79	90.20	91.00
d'Alembert Network	83.19	84.37	83.89	92.23	90.86	91.52

TABLE V
PERFORMANCE OF MODELS ON LOVEDA-FOREST DATASET [24].

Model	IoU (%)		mIoU (%)		Accuracy (%)	
	Forest	Back-ground	Forest	Back-ground	Forest	Back-ground
Deeplabv3+	64.22	75.88	70.05	80.37	84.85	82.61
Segformer	64.63	76.31	70.47	80.36	85.35	82.86
Pidnet-s	64.36	74.08	69.22	84.82	80.85	82.84
Mask2former	65.67	76.83	71.25	81.69	85.30	83.50
Pspnet	62.68	77.08	69.88	73.93	89.18	81.56
Segnext	64.42	76.96	70.69	78.31	87.01	82.66
Knet-s3-r50	65.99	76.16	71.08	84.11	83.45	83.78
SegForest	68.38	79.04	73.71	82.98	87.14	85.06
d'Alembert Network	69.12	80.05	74.34	83.28	87.93	85.87

For the DeepGlobe-Forest dataset, the d'Alembert Network exhibited the highest metrics in all categories: achieving Intersection over Union (IoU) scores of 83.19% for forest and 84.37% for background, mean IoU (mIoU) of 83.89% and 92.23% respectively, and accuracy scores of 90.86% and 91.52%, respectively. These results not only improve upon other advanced models such as SegForest, Pspnet, and Mask2former but also emphasize the network's ability to finely discriminate between forest and non-forest regions, capturing both varying textures of forest landscapes and clear delimitations of background areas. Similarly, on the LoveDA-Forest dataset, the d'Alembert Network again outperformed competing models, recording the highest IoU for the forest at 69.12% and for the background at 80.05%. It also achieved the highest mIoU scores of 74.34% for forest and 83.28% for background, along with accuracy figures of 87.93% and 85.87%, respectively, showing the model's robustness and its enhanced capability in processing and analyzing remote sensing imagery with high precision, particularly in diverse and dynamic environmental settings.

E. Ablation study

The overly small spatial size of the input image patch leads to a significant loss of important information due to an inadequate receptive field. Conversely, an excessively large spatial size of the input image patch introduces many noisy pixels and suffers from inter-class contamination. Therefore, we have established the spatial size of the input image patch

within the range of 5×5, 7×7, 9×9, 11×11, 13×13, 15×15 to evaluate the classification performance across various spatial dimensions. The classification results for two data sets are shown in Figure 3.




Fig. 3. Validation of the optimal hyperparameters with different patch sizes

Line plots demonstrate the trends in training time, accuracy, and loss rate across various batch sizes for the DeepGlobe-Forest and LoveDA-Forest datasets. For both datasets, as the batch size increases from 16 to 128, the training time consistently decreases, indicating enhanced computational efficiency with larger batches, possibly due to fewer updates needed per epoch. This reduction in training time does not seem to compromise the models' ability to learn effectively, as evidenced by the general increase in accuracy with larger batch sizes. However, the most notable improvement is observed in the loss rates, which decrease significantly as the number of batches increases, suggesting that larger batches help the model to converge more smoothly to a lower loss. This trend reflects the trade-off between computational speed and the stability of the training process, where larger batches provide a more stable but potentially less precise gradient estimation, beneficial for the overall learning process of the model.

We have also varied batch sizes to determine their impact on model performance and training dynamics. Batch sizes of 16, 32, 64, and 128 were systematically tested in multiple training

iterations to observe how they influenced the convergence rate, precision and computational efficiency of the neural network. The experiment aimed to identify the optimal batch size that balances between adequate gradient estimation and efficient resource utilization. The results, including metrics such as training time, model accuracy, and loss convergence rates, were carefully recorded and analyzed to deduce the effects of batch size adjustments on the overall effectiveness of the training process. The results are shown in Figure 4




Fig. 4. Validation of the optimal hyperparameters with different batch sizes

Line plots illustrate how different patch sizes affect training time, accuracy, and loss rates for the DeepGlobe-Forest and LoveDA-Forest datasets. As the patch size increases from 5x5 to 15x15, both datasets exhibit a general decrease in training time, suggesting that larger patch sizes enable more efficient training, potentially due to fewer total iterations needed across the dataset. Accuracy trends upward for both datasets as patch size increases, peaking around 11x11 or 13x13 before slightly declining, which indicates an optimal range for capturing relevant features without introducing too much noise or suffering from interclass contamination. Loss rates consistently decrease as the size of the patch increases, reflecting the improved performance of the model with larger patches, which could be attributed to the models' increased ability to capture more comprehensive information about the image, thus potentially

enhancing their learning capability. However, the slight decline in accuracy at the largest patch size suggests a trade-off, where too large a patch might start to incorporate irrelevant information or noise, negatively impacting model precision.

F. Discussion and conclusions

We believe, our approach of using the d'Alembert Convolutional Network in smart forest management, particularly for remote sensing image (RSI) change detection (CD), is a valid alternative compared to existing solutions, for example, the now well-established Spectral-Temporal Transformer (STT) [32]. Both methodologies aim to efficiently capture spectral-temporal features in HSIs, but employ different mechanisms and underlying theories. The STT focus on global spectral-temporal receptive fields with group-wise spectral embedding, linear projection, transformer encoders with an efficient multi-head self-attention mechanism, and a multilayer perceptron head for final change detection. Our approach is different though as it can simultaneously capture spatial features, such as edges and textures, and temporal changes, such as growth patterns or environmental dynamics. The convolution operation in this network is enhanced to include differential operators akin to the d'Alembert operator, creating a more sophisticated mechanism for feature extraction in RSI CD tasks. Furthermore, while STT employs a transformer-based approach with efficient MHSA to reduce computational intensity, the d'Alembert Convolutional Network utilizes the d'Alembert-enhanced kernel in its convolutional layers. This difference in approach leads to a variance in how each model handles the spectral-temporal data, with the d'Alembert network offering a novel perspective by incorporating wave equation principles into image analysis.

To evaluate the efficacy of the proposed d'Alembert convolution approach, we conducted a series of experiments with various parameters on three extensively utilized remote sensing image datasets, designed for change detection. The performance of our method was benchmarked against different established methods. From the analysis of the detection results, it is evident that our d'Alembert approach exhibits good performance in forest change detection, surpassing the comparative methods. For the DeepGlobe-Forest dataset, the d'Alembert Network achieved IoU scores of 83.19% for forest and 84.37% for background, and accuracy scores of 90.86% and 91.52%. These results outperform models like SegForest, Pspnet, and Mask2former benchmarked in the research of Wang et al. [24], showcasing the network's ability to distinguish forest textures and background areas effectively. Similarly, on the LoveDA-Forest dataset, the d'Alembert Network excelled with the highest IoU scores of 69.12% for forest and 80.05% for background, with accuracy rates of 87.93% and 85.87%.

Future work requires further balancing model of complexity and computational resources, considering that while the d'Alembert Convolutional Network enhances spatiotemporal feature extraction, it requires substantial computational power and memory, potentially increasing latency. We're currently working on reducing the number of convolutional layers and

employing techniques like model pruning and quantization to decrease computational load while striving to maintain acceptable performance levels. Naturally, more tests are needed to fully assess the robustness of the model against noise or missing data in the much larger variety of remote sensing images.

FUNDING INFORMATION

This research paper has received funding from Horizon Europe Framework Programme (HORIZON), call Teaming for Excellence (HORIZON-WIDERA-2022-ACCESS-01-two-stage) - Creation of the centre of excellence in smart forestry "Forest 4.0" No. 101059985. This research has been co-funded by the European Union under the project "FOREST 4.0 - Ekscelencijos centras tvariai miško bioekonomikai vystyti" (Nr. 10-042-P-0002)

REFERENCES

- [1] S. Trumbore, P. Brando, and H. Hartmann, "Forest health and global change," *Science*, vol. 349, pp. 814–818, 2015.
- [2] I. Boyd, P. Freer-Smith, C. Gilligan, and H. C. Godfray, "The consequence of tree pests and diseases for ecosystem services," *Science*, vol. 342, p. 1235773, 2013.
- [3] D. Lindenmayer, "Future directions for biodiversity conservation in managed forests: indicator species, impact studies, and monitoring programs," *Forest Ecology and Management*, vol. 115, pp. 277–287, 1999.
- [4] M. B. Nuwantha, C. N. Jayalath, M. P. Rathnayaka, D. C. Fernando, L. Rupasinghe, and M. Chethana, "A drone-based approach for deforestation monitoring," in *2022 13th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 2022.
- [5] W. Zou, W. Jing, G. Chen, Y. Lu, and H. Song, "A survey of big data analytics for smart forestry," *IEEE Access*, vol. 7, pp. 46621–46636, 2019.
- [6] V. Upadhyay and A. Kumar, "Hyperspectral remote sensing of forests: Technological advancements, opportunities and challenges," *Earth Science Informatics*, vol. 11, pp. 487–524, 2018.
- [7] A. Shukla and R. Kot, "An overview of hyperspectral remote sensing and its applications in various disciplines," *IRA-International Journal of Applied Sciences*, vol. 5, pp. 85–90, 2016.
- [8] B. Banerjee, S. Raval, and P. Cullen, "Uav-hyperspectral imaging of spectrally complex environments," *International Journal of Remote Sensing*, vol. 41, pp. 4136–4159, 2020.
- [9] M. Teke, H. S. Deveci, O. Haliloglu, S. Gurbuz, and U. Sakarya, "A short survey of hyperspectral remote sensing applications in agriculture," in *2013 6th International Conference on Recent Advances in Space Technologies (RAST)*, pp. 171–176, 2013.
- [10] R. Prasad and K. Rajan, "Is current forest landscape research approaches providing the right insights? observations from india context," *Ecological Questions*, vol. 20, pp. 85–92, 2015.
- [11] R. E. O. Schultz, T. M. Centeno, G. Selleron, and M. Delgado, "A soft computing-based approach to spatio-temporal prediction," *International Journal of Approximate Reasoning*, vol. 50, pp. 3–20, 2009.
- [12] Y. Chen, H. Jiang, C. Li, X. Jia, and P. Ghamisi, "Deep feature extraction and classification of hyperspectral images based on convolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 10, pp. 6232–6251, 2016.
- [13] S. K. Roy, R. Mondal, M. E. Paoletti, J. M. Haut, and A. Plaza, "Morphological convolutional neural networks for hyperspectral image classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 8689–8702, 2021.
- [14] T. K. Saha, H. Sajjad, Roshani, M. H. Rahaman, and Y. Sharma, "Exploring the impact of land use/land cover changes on the dynamics of deepor wetland (a ramsar site) in assam, india using geospatial techniques and machine learning models," *Modeling Earth Systems and Environment*, 2024.
- [15] B. B. Thien, V. T. Phuong, and D. T. V. Huong, "Detection and assessment of the spatio-temporal land use/cover change in the thai binh province of vietnam's red river delta using remote sensing and gis," *Modeling Earth Systems and Environment*, vol. 9, no. 2, p. 2711 – 2722, 2023.
- [16] R. N. Masolele, V. De Sy, M. Herold, D. Marcos Gonzalez, J. Verbesselt, F. Gieseke, A. G. Mullissa, and C. Martius, "Spatial and temporal deep learning methods for deriving land-use following deforestation: A pan-tropical case study using landsat time series," *Remote Sensing of Environment*, vol. 264, 2021.
- [17] R. V. Mareto, L. M. G. Fonseca, N. Jacobs, T. S. Körting, H. N. Bendini, and L. L. Parente, "Spatio-temporal deep learning approach to map deforestation in amazon rainforest," *IEEE Geoscience and Remote Sensing Letters*, vol. 18, no. 5, p. 771 – 775, 2021.
- [18] R. V. Fonseca, R. G. Negri, A. Pinheiro, and A. M. Atto, "Wavelet spatio-temporal change detection on multitemporal sar images," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, p. 4013 – 4023, 2023.
- [19] M. Dimiyati, D. A. Umarhadi, I. Jamaluddin, D. Awanda, and W. Widyatmanti, "Mangrove monitoring revealed by mdprepost-net using archived landsat imageries," *Remote Sensing Applications: Society and Environment*, vol. 32, 2023.
- [20] W. Jing, T. Lou, Z. Wang, W. Zou, Z. Xu, L. Mohaisen, C. Li, and J. Wang, "A rigorously-incremental spatiotemporal data fusion method for fusing remote sensing images," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, p. 6723 – 6738, 2023.
- [21] I. Demir, K. Koperski, D. Lindenbaum, G. Pang, J. Huang, S. Basu, F. Hughes, D. Tuia, and R. Raskar, "Deepglobe 2018: A challenge to parse the earth through satellite images," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [22] J. Wang, Z. Zheng, A. Ma, X. Lu, and Y. Zhong, "Loveda: A remote sensing land-cover dataset for domain adaptive semantic segmentation," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (J. Vanschoren and S. Yeung, eds.)*, vol. 1, Curran Associates, Inc., 2021.
- [23] J. Wang, Z. Zheng, A. Ma, X. Lu, and Y. Zhong, "LoveDA: A remote sensing land-cover dataset for domain adaptive semantic segmentation," Oct. 2021.
- [24] H. Wang, C. Hu, R. Zhang, and W. Qian, "Segforest: A segmentation model for remote sensing images," *Forests*, vol. 14, p. 1509, July 2023.
- [25] S. C. Yurtkulu, Y. H. Şahin, and G. Unal, "Semantic segmentation with extended deeplabv3 architecture," in *2019 27th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2019.
- [26] J. Xu, Z. Xiong, and S. P. Bhattacharyya, "Pidnet: A real-time semantic segmentation network inspired by pid controllers," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19529–19539, 2023.
- [27] X. Zhu, Z. Cheng, S. Wang, X. Chen, and G. Lu, "Coronary angiography image segmentation based on pspnet," *Computer Methods and Programs in Biomedicine*, vol. 200, p. 105897, 2021.
- [28] W. Zhang, J. Pang, K. Chen, and C. C. Loy, "K-net: Towards unified image segmentation," *Advances in Neural Information Processing Systems*, vol. 34, pp. 10326–10338, 2021.
- [29] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "Segformer: Simple and efficient design for semantic segmentation with transformers," *Advances in neural information processing systems*, vol. 34, pp. 12077–12090, 2021.
- [30] H. Zhang, F. Li, H. Xu, S. Huang, S. Liu, L. M. Ni, and L. Zhang, "Mp-former: Mask-piloted transformer for image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18074–18083, 2023.
- [31] M.-H. Guo, C.-Z. Lu, Q. Hou, Z. Liu, M.-M. Cheng, and S.-M. Hu, "Segnext: Rethinking convolutional attention design for semantic segmentation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 1140–1156, 2022.
- [32] X. Li and J. Ding, "Spectral-temporal transformer for hyperspectral image change detection," *Remote Sensing*, vol. 15, no. 14, 2023.

Optimization of the Cell-based Software Architecture by Applying the Community Detection Approach

Miloš Milić
0000-0002-2521-7607
University of Belgrade
Faculty of Organizational
Sciences, Belgrade
Jove Ilića 154,
11000 Belgrade, Serbia
Email: milos.milic@fon.bg.ac.rs

Dragana Makajić-Nikolić
0000-0002-0790-6791
University of Belgrade
Faculty of Organizational
Sciences, Belgrade
Jove Ilića 154,
11000 Belgrade, Serbia
Email: dragana.makajic-
nikolic@fon.bg.ac.rs

Abstract—The aim of this research is to present the Cell-based software architecture and explore its optimization. Cell-based software architecture organizes a software system into interconnected cells, each containing multiple elements. This research focuses on optimizing cell-based architecture, particularly the number of cells and their internal organization. In this context, the Community Detection approach, which identifies closely connected elements, was applied. Additionally, the model incorporates the concept of functionality, defined as a set of capabilities allowable and actionable by the software system. We conducted a series of experiments based on the defined mathematical model to validate our approach, achieving optimal and near-optimal solutions within a given time limit. Considering that each cell can contain multiple elements realized in various architectural styles, the proposed model allows for the integration of different architectures within the same software system. This flexibility enhances the system's overall adaptability and efficiency.

Index Terms—software architecture, cell-based architecture, community detection, architecture optimization.

I. INTRODUCTION

IN TODAY'S digital age, the application of software systems spans across various domains. These systems enable seamless communication and data exchange within and across different industries. In the interconnected world, software systems can be utilized by a diverse range of clients, and it is essential to ensure they have capabilities to support them effectively.

In addition to functional requirements, these capabilities are related to non-functional requirements such as security, deployability, availability, scalability, reliability, resilience, maintainability, etc. [1]. However, achieving a high level of non-functional requirements can be a challenging task. Non-functional requirements are typically defined as quality at-

tributes of a software system, and are closely related to software architecture [2].

Software architecture of a system can be defined as the set of elements needed to reason about the system [3], encompassing various software components, their relationships, as well as the properties of components and relationships [1]. Software architecture can be considered as a blueprint for further software design, based on which various components are created. In this context, software architects and engineers should consider software architecture from the earliest phases of development.

This research presents the Cell-based software architecture. Cell-based architecture considers the organization of a software system in the form of interconnected cells, while each cell can contain multiple elements [4]-[5]. The research observes cell-based software architecture optimization, specifically focusing on the number of cells and their internal organization, as community detection problem.

The rest of the paper is organized as follows. Section 2 introduces various software architectures that can be applied in the software development process. Additionally, the Cell-based software architecture is presented, as well as the Community Detection problem and its application in different fields. Section 3 introduces the problem and defines a mathematical model for Cell-based software architecture optimization. Evaluation and optimization results are presented in Section 4. Finally, the conclusion is presented in Section 5.

II. BACKGROUND

This section introduces various software architectures, with a focus on Cell-based software architecture. Given that cells can be represented as a network of connected nodes, the section also covers the Community Detection problem and its application in various fields.

This work was supported by the University of Belgrade – Faculty of Organizational Sciences.

A. Software Architecture

When software design is concerned, various software architectures can be observed. For example, monolithic architecture represents a traditional software design approach. This architecture involves multiple modules that are executed together as a single unit at runtime, resulting in high coupling between the modules [6]. On the other hand, microservice architecture is an alternative to monolithic architecture. In microservice architecture, each element is implemented as a separate microservice, which operates independently as a single unit at runtime. This approach results in low coupling between microservices [6]. However, taking into account that each microservice is managed independently, microservice organization and communication must be carefully considered [7]. In addition, microservices typically require additional components for management, such as microservice orchestration and choreography [8], which can introduce additional complexity. Although monolith and microservice architecture can co-exist within the same system, researchers are exploring approaches decomposing and gradually transitioning from monolithic applications to microservices [9]-[12]. Both monolithic and microservice architectures require infrastructure services (e.g., application server, database server, etc.), which can be either on-premises or cloud-based.

Another alternative to monolithic and microservice architectures is serverless software architecture, an approach that focuses on designing services related to specific business capabilities [13]. In this context, Functions-as-a-Service (FaaS) can be coded and deployed, while the underlying infrastructure is managed by the cloud provider [14]. Although this approach allows software engineers to focus on business functions, it results in a high degree of coupling with the infrastructure services provisioned by the cloud provider.

Based on the previous discussion, it can be stated that each software architecture has its own pros and cons that should be carefully considered during the software design process.

B. Cell-based Software Architecture

Cell-based architecture can be defined as a software architecture that incorporates multiple units of workload, with each unit known as a cell [5]. Each cell is independent from other cells, does not share state with other cells, and can encapsulate multiple components of different types [4]-[5]. Additionally, each cell contains a cell gateway, serving as the central entry point for cell communication. In this context, intra-cell and inter-cell communication can be observed, which is realized with well-defined interfaces and protocols [4]. A specific set of functionalities or services can be incorporated within a cell, defining a cell boundary. In this context, cell-based architecture can be related with domain-driven software design [15].

Conceptual overview of the Cell-based software architecture is presented in Figure 1. The figure depicts two cells with multiple elements, with the cell boundaries outlined by octagons. Cell A incorporates three elements (e.g., one monolith and two microservices), while Cell B also includes three elements (e.g., three microservices). Additionally, element A2

from Cell A communicates with Cell B through the cell gateway. In this way inter-cell communication is realized [4]. On the other hand, element B2 communicates with element B3. This communication is performed inside Cell B and represents intra-cell communication [4]. Each cell is autonomous and can be managed independently of other cells. As a result, better encapsulation, isolation, and distribution of software architecture elements can be achieved, addressing some of the typical challenges in software architecture design [16].




Fig 1. Conceptual overview of the Cell-based software architecture

Considering that a cell can incorporate multiple capabilities implemented in various architectures, the cell-based approach facilitates the introduction of multi-architecture software development. In this context, the benefits of each applied architecture can be utilized, while their cons can be managed. This approach allows each cell to be independent and iterate individually, resulting in decentralized software architecture [4].

C. Community Detection Problem

Community detection problem belongs to the field of Complex Network Analysis. Its most common areas of application are: social networks [17]-[18], neuroscience and biology [19], supply chain networks [20]-[21], politics, customer segmentation, smart advertising and targeted marketing [22], etc. Community detection approach were also applied in software engineering since the process-oriented and the object-oriented software architecture both can be presented as complex network [23] characterized by properties like those commonly observed in other complex networks [24]. Authors Pan, Jing, and Li used community detection approach for refactoring the package structures of object-oriented software in order to improve the maintenance process [25]. Software maintenance was also emphasized as the reason for using community detection in research conducted by Huang et al. [26]. Authors Hou, Yao, and Gong applied community detection approach to developer collaboration network in software ecosystem based on developer cooperation intensity [27].

Communities are groups of network's vertices with the common properties and/or role in the network [28]. The community detection problem is to find communities that maximize a given quality function. The solution of the problem is a set of communities such that the number of edges within the

community is greater of the number of edges between the community's vertices and the rest of the network (Figure 2).




Fig 2. A simple graph with three communities

There are several quality measures intended to evaluate the structure of communities [29]. In this paper we use the standard and most used measure of quality, called Newman–Girvan modularity [30]. One of the formulations of Newman–Girvan modularity is:

$$Q = \sum_{m \in C} \left(\frac{L_m}{L} - \left(\frac{D_m}{2L} \right)^2 \right) \quad (1)$$

where C represents the set of communities, L_m is the sum of the weights of the edges within community m , L is the sum of weights of all edges in entire the network, and D_m is the degree of the vertices in community m .

The function (1) is nonlinear, and it can only be solved for small and medium-sized unweighted graphs [31]. Hence, several solving methods and linearizations of it can be found in the literature [32]–[33]. In this paper, we use the variant of modularity function (1) from [34] that enables further linearization:

$$Q = \sum_{k \in C} \left(\frac{1}{L} \sum_{(i,j) \in E_m} b_{ij} - \frac{1}{4L} \sum_{(i,j) \in E_m} d_i d_j \right) \quad (2)$$

where E_m represents the set of edges in community m of a given graph $G=(E,V)$, V is the set of the vertices. Parameter b_{ij} is the weight of the edge (i,j) , $(i,j) \in E$, and d_i is the weight of the vertex, obtained as sum of weight of all input and output edges of vertex i :

$$d_i = \sum_{(i,j),(j,i) \in E} b_{ij}, \quad i \in V \quad (3)$$

III. MATHEMATICAL MODEL FOR OPTIMIZING CELL-BASED SOFTWARE ARCHITECTURE

As previously discussed in Section 2, cell-based software architecture can be depicted as a network of interconnected cells and elements that communicate with each other. Given that the solution of Community Detection problem can identify closely connected items, this section presents a mathematical model for optimizing cell-based software architecture. Optimizing the cell-based architecture can potentially lead to

better resource utilization through an optimal number of cells and their internal organization. Additionally, various software quality attributes can be enhanced.

When the software architecture whose elements should be grouped into cells based on community detection problem, the elements of the architecture are vertices of the graph $G=(E,V)$. The edges of the graph exist between the elements (vertices) which communicate, while the weight of the edge (i,j) , b_{ij} represents the intensity of the communication.

Since the weights inside the parentheses in equation (2) should be calculated only for the edges and vertices belonging to the same community, the set of communities C and binary variables y_{ik} are introduced:

$$y_{ik} = \begin{cases} 1 & \text{if } i\text{-th vertex is in community } k, \\ 0 & \text{otherwise} \end{cases}, \quad i \in V, k \in C.$$

Equation (2) now became:

$$Q = \sum_{k \in C} (L_1 \sum_{(i,j) \in E} b_{ij} y_{ik} y_{jk} - L_2 \sum_{(i,j) \in E} d_i d_j y_{ik} y_{jk}) \quad (4)$$

where $L_1 = 1/L$, $L_2 = 1/4L$, and L is the sum of weights of all edges in entire the network.

The nonlinearity $y_{ik} y_{jk}$ could be replaced by auxiliary binary variables z_{ijk} :

$$z_{ijk} = \begin{cases} 1 & \text{if edge } (i,j) \text{ is in community } k, \\ 0 & \text{otherwise} \end{cases}, \quad (i,j) \in E, k \in C$$

and inequalities:

$$z_{ijk} \geq y_{ik} + y_{jk} - 1, \quad k \in C, (i,j) \in E \quad (5)$$

$$z_{ijk} \leq y_{ik}, \quad k \in C, (i,j) \in E \quad (6)$$

$$z_{ijk} \leq y_{jk}, \quad k \in C, (i,j) \in E$$

The condition (5) ensures that variable z_{ijk} get the value 1 if both y_{ik} and y_{jk} have the value 1, i.e. the edge (i,j) is inside the community k if both vertices i and j belong to the community k . Since based on condition (5), value of z_{ijk} can be 1 if y_{ik} and/or y_{jk} are equal to zero, the condition (6) is introduced to prevent such solutions. Furthermore, if for an edge (i,j) z_{ijk} equals zero for all $k \in C$, it indicates that edge (i,j) does not belong to any community; instead, it represents a link between two different communities.

Additionally, the model also incorporates the concept of functionality. Functionality can be defined as a set of capabilities allowable and actionable by the software system [35]. Each functionality contains elements focused on a specific domain and should not be mixed to maintain boundaries, reduce complexity, and ensure modularity. In a cell-based software architecture, a single functionality can be represented by one or more cells, forming the foundation for optimizing the software architecture. In addition, different functionalities should not be organized in the same cell, allowing better separation of concerns between cells. As a result, each cell can be independent and managed individually [4].

In addition to the already introduced parameters and variables, notation used for the mathematical model formulation is as follows.

Sets:

- FC - set of functionalities,
- F_l - set of l -th functionality,
- $F_l \subset FC, \bigcap_{l \in FC} F_l = \emptyset, \bigcup_{l \in FC} F_l = FC$

Parameters:

- e - lower bound of the number of vertices in communities,

Variables:

$$x_k = \begin{cases} 1 & \text{if } k\text{-th community exist} \\ 0 & \text{otherwise} \end{cases}, k \in C$$

The proposed mathematical model is listed below.

$$\max f(z) = \sum_{k \in C} (L_1 \sum_{(i,j) \in E} b_{ij} z_{ijk} - L_2 \sum_{(i,j) \in E} d_i d_j z_{ijk}) \quad (7)$$

s.t.

$$z_{ijk} \geq y_{ik} + y_{jk} - 1, k \in C, (i, j) \in E \quad (8)$$

$$z_{ijk} \leq y_{ik}, k \in C, (i, j) \in E \quad (9)$$

$$\sum_{k \in C} y_{ik} = 1, i \in V \quad (10)$$

$$y_{ik} \leq x_k, k \in C, i \in V \quad (11)$$

$$\sum_{i \in V} y_{ik} \geq e \cdot x_k, k \in C \quad (12)$$

$$y_{ik} + y_{jk} \leq 1, k \in C, i \in F(l), j \in F(p), l, p \in FC, l \neq p \quad (13)$$

$$x_k \in \{0, 1\}, k \in C \quad (14)$$

$$y_{ik} \in \{0, 1\}, i \in V, k \in C \quad (15)$$

$$z_{ijk} \in \{0, 1\}, (i, j) \in E, k \in C \quad (16)$$

The objective function (7) represents the modularity measure linearized by replacing $y_{ik}y_{jk}$ with z_{ijk} in (4). Since this function should be maximized, the first addend in parentheses will be as large as possible. Thus, the branches that have a greater weight will be within the same community, that is, the software elements with more frequent communication will be in the same cell. Constraints (8-9) are related to linearization. Constraint (10) ensures that each vertex is assigned exactly to one community. The constraint (11), the value 1 is set to x_k if some vertex is assigned to the k -th community. Constraint (12) is related to the minimal number of vertices assigned to the existing communities. Constraint (13) provides that vertices of different functionality cannot belong to the same community, i.e. only vertices of the same functionality can be in the same community. Constraints (14-16) are related to the binary restrictions on the variables.

If necessary, additional constraints can be introduced. For example, although the parameter e defines the lower bound for the number of vertices in communities, an additional constraint can be added to specify a different minimum number of vertices for a particular community. This allows for fine-grained definition of cell structure in specific circumstances.

For example, if some of the software element should be isolated in a cell, a set of such element $VI \subset V$ and additional constraint can be included into mathematical model:

$$y_{ik} + y_{jk} \leq 1, k \in C, i \in VI, j \in V, j \neq i \quad (17)$$

Additionally, if some elements should be in the same cell, regardless the connections between them, the mathematical model can be extended as follows.

- G - set of predefined groups of elements,
- V_q - set of elements predetermined to be in the same cell, $q \in G$,

$$y_{ik} = y_{jk}, k \in C, i, j \in V_q, q \in G \quad (18)$$

$$y_{ik} + y_{jk} \leq 1, k \in C, i \in V_q, j \in V \setminus V_q \quad (19)$$

The constraint (18) ensures that the predetermined elements are in the same cell but allows other elements to be assigned to that cell as well. If it is necessary to assign to the same cell only the elements from $q \in G$, constraint (19) should be included into mathematical model.

IV. EVALUATION

The optimized structure of the solution can be graphically presented based on the results.

Figure 3 presents the results of the optimization of a manufacturing software system (e.g., mobile phone manufacturing). The input includes defined Production and Purchasing functionalities. The Production functionality comprises two monolithic applications (i.e., *production* and *legacy*) and two microservices (i.e., *inventory* and *product*), while the Purchasing functionality consists of three microservices (i.e., *order*, *payment*, and *notification*). Additionally, the communication between these elements is specified (the weights of the edges in Figure 3).

Based on the performed optimization, the resulted solution includes three communities (named Community A, Community B, and Community C), each containing different elements (see Figure 3). In the following text these communities will be referred to as the Production Cell (Community A), Purchasing Cell (Community B), and Legacy Cell (Community C).

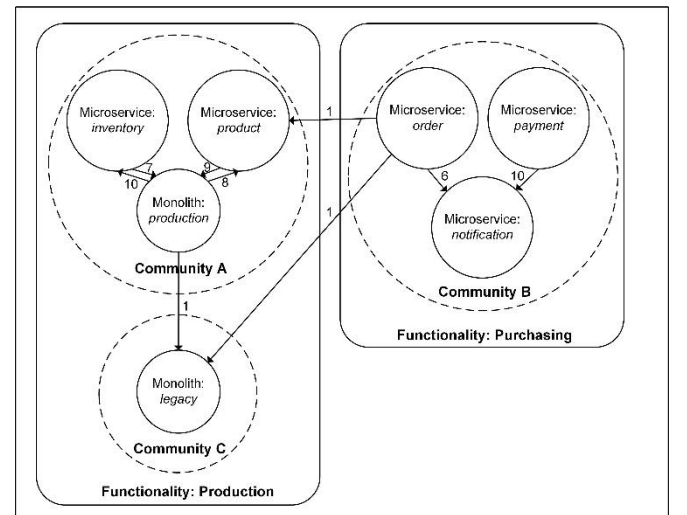


Fig 3. Results of the optimization of a manufacturing software system

An additional observation pertains to the Legacy Community, which contains only one element (i.e., the *legacy* monolith). A legacy software system is defined as a core system that has been functioning correctly in production for decades [36]. Considering the prevalence of legacy systems today, researchers are exploring approaches to migrate these systems to modern architectures [37]-[39]. In the context of cell-based software architecture, the legacy element is incorporated within a specific cell. From the optimization model perspective, this is represented as an additional constraint that restricts the particular cell structure:

$$y_{legk} + y_{jk} \leq 1, k \in C, j \in V \setminus leg \quad (20)$$

where *leg* is the index of the variable corresponding to the *legacy* monolith. This constraint ensures that no other element can be in the cell containing the *legacy* monolith.

Another interesting observation pertains to the Cell Gateway component. While this component is not explicitly represented in the mathematical model, it serves as the central entry point for cell communication [4]. Therefore, we have included one gateway per cell based on the optimal solution. The final software architecture is shown in Figure 4.

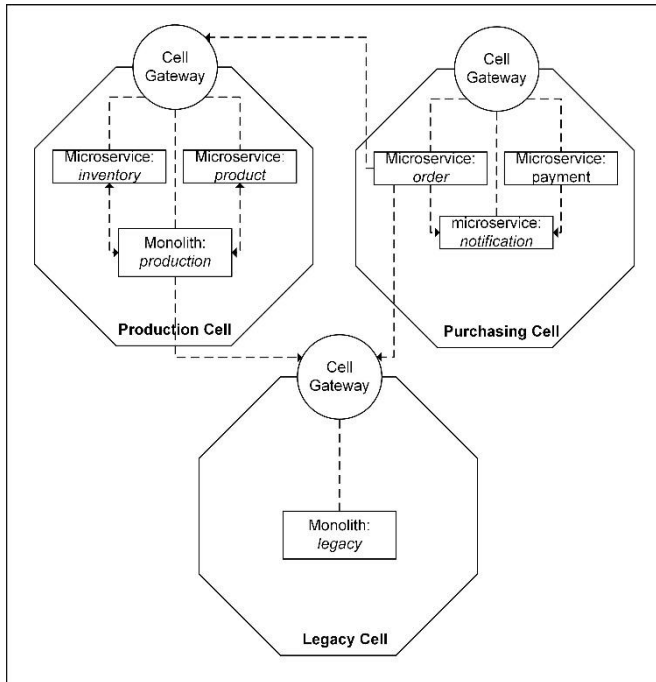


Fig 4. Cell-based software architecture based on the optimal solution

In order to validate the mathematical model and examine the dimensions of the problem that can be solved exactly, we conducted a series of experiments.

The evaluation considered three functionalities, with the number of elements varied. Each element represents an instance of software architecture encapsulating specific features (e.g., a monolith with multiple features, a microservice containing a single feature). Considering that these elements can vary in terms of their applied software architecture and size, a software system can encompass numerous elements. Within

this context, the first functionality incorporated 40% of the elements, while the remaining elements were equally divided between the other two functionalities. Taking into account that elements cooperate with each other, each element has at least one connection with another element within the same functionality, while up to 40% of elements have double connections within the same functionality. Finally, considering that all functionalities are part of the same software system, two connections between elements from different functionalities are also established. The previously discussed elements, such as edges and their weights (the intensity of the communication between the elements - parameter b_{ij}) are randomly generated. The lower bound of the number of vertices in communities (parameter e) is set to 2.

The experiments were conducted for the graphs whose dimensions are given in Table 1. The columns named vertices and edges give the number of vertices and edges, respectively, while column L represents the total weight of all vertices in graphs.

TABLE I.
A SUMMARY OF DATA SETS USED IN THE EXPERIMENTS

Case	Number of vertices	Number of edges	L
1	30	38	163
2	40	52	214
3	50	64	252
4	60	77	307
5	70	90	351
6	80	103	393
7	90	116	462
8	100	128	519
9	150	193	746
10	200	257	991

Given that the model parameters are randomly generated, the mathematical model presented in section III was applied on ten instances of each graph from Table 1. All optimizations were performed solved using GLPK software on a laptop computer equipped with 11th Gen Intel(R) Core(TM) i5 and 16 GB of RAM. The solving method used in GLPK software was Branch and Cut, with Gomory's mixed integer cuts, MIR (mixed integer rounding) cuts, mixed cover cuts and clique cuts options. Execution time was limited, depending on graphs dimensions.

Time limitations were: 5 minutes for cases 1 to 3, 10 minutes for cases 4 to 7, and 15 minutes for cases 8 to 10. Table 2 gives the optimization time. The second column gives the number of instances (out of ten) for which the optimal solution was found within the given time limit. The three right columns give the minimum, maximum, and average duration (in seconds) of the optimization among ten generated instances.

TABLE II.
DURATIONS OF THE SUCCESSFUL OPTIMIZATIONS

Case	Number of optimal solutions	min	max	avg
1	8	0.2	12.9	1.9
2	9	0.8	8.1	1.8
3	9	1.7	86.7	12.0
4	8	6	13.9	8.7
5	8	11.2	181.7	39.9
6	8	20.4	61	32.5
7	6	30.4	58.1	44.6
8	7	30.2	44.2	39.2
9	7	87.3	218.8	117.3
10	5	142	552	423.3

As expected, the number of instances that can be solved within the time limit decreases and optimization time increases with increasing of graphs dimensions. However, even for lower dimensions, the graph topology and parameter values can prevent finding the optimal solution in a given time limit, as can be seen in the case of 30 nodes. Generally, in most instances of graphs of the same dimensions, the optimization times were similar. Figure 5 shows distribution of the optimization time for all solved instances.




Fig 5. Duration of optimizations of all instances of all ten cases

In most cases, the optimization times are grouped, except for cases 3, 5, and 9 where the optimization of one instance takes significantly longer than the others and in case 10 where the optimization of one instance is significantly faster than the others. The reason for these deviations lies in the topology of the graphs and the values of the model parameters.

Table 3 shows the average performance of obtained solutions. The last column represents the optimal value of the objective function, i.e. linearized modularity function.

TABLE III.
AVERAGE PERFORMANCE OF THE OPTIMAL SOLUTIONS

Case	Number of cells	Min cell size	Max cell size	Modularity value
1	8.63	2	5.13	4.25E-07
2	10.00	2	7.10	-2.78E-03
3	10.78	2	8.11	-5.88E-02
4	10	2	9.57	1.75E-07
5	11.63	2.25	11.5	1.89E-06
6	13	2.5	13.25	1.44E-06
7	11	2.2	16.67	-1.61E-02
8	10.14	2.43	17.43	8.86E-07
9	11.71	3.43	25.43	1.14E-06
10	15.80	2	32.00	0

The number of cells increases slowly with the increasing of graphs dimensions even though the model does not contain its upper limit. Additionally, almost 20% of cells consist of two elements in the majority of instances. The only exception among 75 successfully solved instances is the two instances of the graphs with 70, 80, 100 and 150 vertices and one instance of the graph with 90 vertices. Figure 6 shows the number of cells for all solved instances.




Fig 6. The number of cells in all instances of all ten cases

Based on Figure 6, it can be concluded that the number of cells is generally grouped for most instances within the same case. However, the number of cells slightly varies with different problem dimensions and typically ranges from 7 to 16, with a few exceptions. This indicates that the number of cells is more influenced by the topology of the correspondence graph and the parameters values than by the number of software elements.

The values of linearized modularity function in Table 3 are small, even negative in some cases. These negative values indicate that nodes are less connected within communities. However, the interpretation of the value of the modularity function in these experiments is not of great importance, given that hypothetical examples with randomly generated graph branches and mathematical model parameters were used.

V. CONCLUSION

The selection of software architecture guides the software design and development process, making it an important topic in the field of software engineering. Therefore, the chosen software architecture should be carefully selected to suit the specific needs of the software system being implemented. While this research focuses on optimizing cell-based software architecture, particularly the number of cells and their internal organization, additional research directions can be considered.

Further research could examine system workload in the context of additional non-functional attributes such as scalability, availability, and reliability. Additionally, the details of intra-cell and inter-cell communication could be further investigated. Regarding mathematical model, different linearization of modularity function can be examined as well as different quality measures. The main goal of this research was to investigate the validity of modeling cell-based software architecture as community detection problem. Since this problem is NP hard, the next step of the research will be to develop a heuristic for solving large scale problems.

DATA AVAILABILITY

Input data and optimization results from a series of experiments can be accessed at the following address: <https://github.com/mmilicfon/fedcsis2024>.

REFERENCES

- [1] P. Bourque and R. E. Fairley (Eds), *Guide to the Software Engineering Body of Knowledge (SWEBOK (R)): Version 3.0*, IEEE Computer Society Press, 2014.
- [2] A. Chandrasekar, S. Rajesh, and P. Rajesh, "A research study on software quality attributes", *International Journal of Scientific and Research Publications*, Vol. 4, No. 1, pp. 14-19, 2014.
- [3] L. Bass, P. Clements, and R. Kazman, *Software Architecture in Practice*, 4th Edition, Addison-Wesley Professional, 2021.
- [4] A. Abeyasinghe, "Cell-Based Architecture: A Decentralized Reference Architecture for Cloud-native Applications", 2024. Available online: <https://github.com/wso2/reference-architecture/blob/master/reference-architecture-cell-based.md> (Access Date: May 17, 2024).
- [5] Amazon Web Services, "Reducing the Scope of Impact with Cell-Based Architecture: AWS Well-Architected", Amazon Corporation, 2024. Available online: <https://docs.aws.amazon.com/wellarchitected/latest/reducing-scope-of-impact-with-cell-based-architecture/reducing-scope-of-impact-with-cell-based-architecture.html> (Access Date: July 12, 2024).
- [6] G. Blinowski, A. Ojdowska, and A. Przybyłek, "Monolithic vs. micro-service architecture: A performance and scalability evaluation", *IEEE Access*, Vol. 10, pp. 20357-20374, 2022, <https://doi.org/10.1109/ACCESS.2022.3152803>.
- [7] S. Hassan and R. Bahsoon, "Microservices and their design trade-offs: A self-adaptive roadmap", in *Proceedings of the 2016 IEEE International Conference on Services Computing (SCC)*, pp. 813-818, IEEE, 2016, <https://doi.org/10.1109/SCC.2016.113>.
- [8] N. Singhal, U. Sakthivel, and P. Raj, "Selection mechanism of micro-services orchestration vs. choreography". *International Journal of Web & Semantic Technology (IJWesT)*, Vol. 10, No. 1, pp. 1-13, 2019, <https://doi.org/10.5121/ijwest.2019.10101>.
- [9] Y. Abgaz, A. McCarren, P. Elger, D. Solan, N. Lapuz, M. Bivol, G. Jackson, M. Yilmaz, J. Buckley, and P. Clarke, "Decomposition of monolith applications into microservices architectures: A systematic review", *IEEE Transactions on Software Engineering*, Vol. 49, No. 8, pp. 4213-4242, 2023, <https://doi.org/10.1109/TSE.2023.3287297>.
- [10] R. Chen, S. Li, and Z. Li, "From monolith to microservices: A dataflow-driven approach", In J. Lv, H. Zhang, X. Liu, and M. Hinchey (Eds.), *Proceedings of the 2017 24th Asia-Pacific Software Engineering Conference (APSEC)*, pp. 466-475, IEEE, 2017, <https://doi.org/10.1109/APSEC.2017.53>.
- [11] G. Mazlami, J. Cito, and P. Leitner, "Extraction of microservices from monolithic software architectures", in *Proceedings of the 2017 IEEE International Conference on Web Services (ICWS)*, pp. 524-531, IEEE, 2017, <https://doi.org/10.1109/ICWS.2017.61>.
- [12] K. Sellami, M. A. Saied, A. Ouni, and R. Abdalkareem, "Combining static and dynamic analysis to decompose monolithic application into microservices", in *International Conference on Service-Oriented Computing*, pp. 203-218, Cham: Springer Nature Switzerland, 2022, https://doi.org/10.1007/978-3-031-20984-0_14.
- [13] M. Sewak and S. Singh, "Winning in the era of serverless computing and function as a service", in *Proceedings of the 2018 3rd International Conference for Convergence in Technology (I2CT)*, pp. 1-5, IEEE, 2018, <https://doi.org/10.1109/I2CT.2018.8529465>.
- [14] M. Shahrad, R. Fonseca, I. Goiri, G. Chaudhry, P. Batum, J. Cooke, E. Laureano, C. Tresness, M. Russinovich, and R. Bianchini, (2020), "Serverless in the wild: Characterizing and optimizing the serverless workload at a large cloud provider", in *Proceedings of the 2020 USENIX annual technical conference (USENIX ATC 20)*, pp. 205-218, 2020.
- [15] S. Millett and N. Tune, *Patterns, principles, and practices of domain-driven design*, John Wiley & Sons, 2015.
- [16] A. Bierska, B. Buhnova, and H. Bangui, "An Integrated Checklist for Architecture Design of Critical Software Systems", *Annals of Computer Science and Information Systems (FedCSIS)*, pp. 133-140, 2022, IEEE, <https://doi.org/10.15439/2022F287>.
- [17] C. C. Lin, J. R. Kang, and J. Y. Chen, "An integer programming approach and visual analysis for detecting hierarchical community structures in social networks", *Information Sciences*, 299, pp. 296-311, 2015, <https://doi.org/10.1016/j.ins.2014.12.009>.
- [18] A. R. Costa, & C. G. Ralha, "AC2CD: An actor-critic architecture for community detection in dynamic social networks", *Knowledge-Based Systems*, 261, 110202, 2023, <https://doi.org/10.1016/j.knosys.2022.110202>.
- [19] E. M. Mohamed, T. Agouti, A. Tikniouine, and M. El Adnani, "A comprehensive literature review on community detection: Approaches and applications", *Procedia Computer Science*, 151, pp. 295-302, 2019, <https://doi.org/10.1016/j.procs.2019.04.042>.
- [20] N. M. Viljoen and J. W. Joubert, "Supply chain micro-communities in urban areas", *Journal of Transport Geography*, 74, 211-222, 2019, <https://doi.org/10.1016/j.jtrangeo.2018.11.011>.
- [21] Z. Lu and Z. Dong, "A Gravitation-Based Hierarchical Community Detection Algorithm for Structuring Supply Chain Network", *International Journal of Computational Intelligence Systems*, Vol. 16, No. 1, 110, 2023, <https://doi.org/10.1007/s44196-023-00290-x>.
- [22] A. Karataş, and S. Şahin, "Application areas of community detection: A review", in *Proceedings of the 2018 International congress on big data, deep learning and fighting cyber terrorism (IBIGDELFT)*, pp. 65-70, IEEE, 2018, <https://doi.org/10.1109/IBIGDELFT.2018.8625349>.
- [23] D. Li, Y. Han, and J. Hu, "Complex network thinking in software engineering", in *Proceedings of the 2008 International Conference on Computer Science and Software Engineering*, pp. 264-268, IEEE, 2008, <https://doi.org/10.1109/CSSE.2008.689>.
- [24] L. Šubelj and M. Bajec, "Community structure of complex software systems: Analysis and applications", *Physica A: Statistical Mechanics and its Applications*, Vol. 390, No. 16, pp. 2968-2975, 2011, <https://doi.org/10.1016/j.physa.2011.03.036>.
- [25] W. F. Pan, B. Jiang, and B. Li, "Refactoring software packages via community detection in complex software networks", *International Journal of Automation and Computing*, Vol.10, No. 2, pp. 157-166, 2013, <https://doi.org/10.1007/s11633-013-0708-y>.
- [26] G. Huang, P. Zhang, B. Zhang, T. Yin, and J. Ren, "The optimal community detection of software based on complex networks", *International Journal of Modern Physics C*, Vol. 27, No. 08, 1650085, 2016, <https://doi.org/10.1142/S0129183116500856>.
- [27] T. Hou, X. Yao, and D. Gong, "Community detection in software ecosystem by comprehensively evaluating developer cooperation intensity", *Information and Software Technology*, 130, 106451, 2021, <https://doi.org/10.1016/j.infsof.2020.106451>.

- [28] S. Fortunato, "Community detection in graphs", *Physics reports*, Vol. 486, No. 3-5, pp. 75-174, 2010, <https://doi.org/10.1016/j.physrep.2009.11.002>.
- [29] V. L. Dao, C. Bothorel, and P. Lenca, "Community structure: A comparative evaluation of community detection methods", *Network Science*, Vol. 8, No. 1, pp. 1-41, 2020, <https://doi.org/10.1017/nws.2019.59>.
- [30] M. E. Newman and M. Girvan, "Finding and evaluating community structure in networks", *Physical review E*, Vol. 69, No. 2, 026113, 2004, <https://doi.org/10.1103/PhysRevE.69.026113>.
- [31] L. Bennett, S. Liu, L.G. Papageorgiou, and S. Tsoka, "A mathematical programming approach to community structure detection in complex networks", *Computer Aided Chemical Engineering*, Vol. 30, pp. 1387-1391, 2012, <https://doi.org/10.1016/B978-0-444-59520-1.50136-6>.
- [32] B. Serrano and T. Vidal, "Community detection in the stochastic block model by mixed integer programming", *Pattern Recognition*, Vol. 152, 110487, 2024, <https://doi.org/10.1016/j.patcog.2024.110487>.
- [33] A. Ferdowsi and M. D. Chenary, "Toward an Optimal Solution to the Network Partitioning Problem", *Annals of Computer Science and Information Systems*, Vol. 35, 18th Conference on Computer Science and Intelligence Systems (FedCSIS), pp. 111-117, 2023, IEEE, <https://doi.org/10.15439/2023F2832>.
- [34] E. Alinezhad, B. Teimourpour, M.M. Sepehri, and M. Kargari, "Community detection in attributed networks considering both structural and attribute similarities: two mathematical programming approaches", *Neural Computing and Applications*, Vol. 32, pp. 3203-3220, 2020, <https://doi.org/10.1007/s00521-019-04064-5>.
- [35] ISO/IEC/IEEE 24765:2017 Systems and software engineering — Vocabulary, International Organization for Standardization, Available online: <https://www.iso.org> (Access Date: May 27, 2024).
- [36] R. Khadka, B. V. Batlajery, A. M. Saeidi, S. Jansen, and J. Hage, "How do professionals perceive legacy systems and software modernization?", in *Proceedings of the ACM 36th International Conference on Software Engineering (ICSE 2014)*, ACM, pp. 36-47, 2014, <http://dx.doi.org/10.1145/2568225.2568318>.
- [37] J. Kazanavičius and D. Mažeika, "Migrating legacy software to microservices architecture", in *Proceedings of the IEEE 2019 Open Conference of Electrical, Electronic and Information Sciences (eStream)*, pp. 1-5, 2019, IEEE, <https://doi.org/10.1109/eStream.2019.8732170>.
- [38] A. Ahmad and M. A. Babar, "A framework for architecture-driven migration of legacy systems to cloud-enabled software", in *Proceedings of the WICSA 2014 Companion Volume*, pp. 1-8, 2014, <http://dx.doi.org/10.1145/2578128.2578232>.
- [39] A. Menychtas, C. Santzaridou, G. Kousiouris, T. Varvarigou, L. Orue-Echevarria, J. Alonso, J. Gorronogoitia, H. Bruneliere, O. Strauss, T. Senkova, B. Pellens, and P. Stuer, "ARTIST Methodology and Framework: A novel approach for the migration of legacy software on the Cloud", in *Proceedings of the IEEE 2013 15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, pp. 424-431, 2013, IEEE, <https://doi.org/10.1109/SYNASC.2013.62>.

Model-Agnostic Machine Learning Model Updating – A Case Study on a real-world Application

Julia Poray

0009-0003-4868-1255

GlobalFoundries Dresden Module One LLC & Co. KG

julia.poray@globalfoundries.com

Bogdan Franczyk*

0000-0002-5740-2946

Leipzig University, Information Systems Institute

Grimmaische Straße 12, 04109 Leipzig

franczyk@wifa.uni-leipzig.de

Thomas Heller*

0009-0006-9221-3784

GlobalFoundries Dresden Module One LLC & Co. KG

thomas.heller@globalfoundries.com

Abstract—The application of developments in the real world is the final aim of all scientific works. In the case of Data Science and Machine Learning, this means there are additional tasks to care about, compared to the rather academic part of “just” building a model based on the available data. In the well accepted Cross Industry Standard for Data Mining (CRISP-DM), one of these tasks is the maintenance of the deployed application. This task can be of extreme importance, since in real-world applications the model performance often decreases over time, usually due to Concept Drift. This directly leads to the need to adapt/update the used Machine Learning model. In this work, available model-agnostic model update methods are evaluated on a real-world industry application, here Virtual Metrology in semiconductor fabrication. The results show that for the real-world use case sliding window techniques performed best. The models used in the experiments were an XGBoost and Neural Network. For the Neural Network, Model-Agnostic Meta-Learning and Learning to learn by Gradient Descent by Gradient Descent were applied as update techniques (among others) and did not show any improvement compared to the baseline of not updating the Neural Network. The implementation of the update techniques was validated on an artificial use case for which they worked well.

Index Terms—Machine Learning, Concept Drift Adaptation, Model Updating, Supervised Regression, Model degradation, Semiconductor Fabrication, Virtual Metrology

I. INTRODUCTION

IN modern industrial environments, e.g. semiconductor fabrication, there is an increasing use of Machine Learning models throughout the whole processing chain and beyond. A high amount of use potential lies in Machine Learning regression tasks which aim to provide hard to obtain quality parameters. One such case is Virtual Metrology [1][2], where measurement data are predicted from process machine information, i.e. sensor data and more. An example for such a measurement is the thickness of a layer deposited on a wafer. Data from such an example are used in this work.

However, there are numerous other use cases as well. Given the dynamic nature of the industrial environment, a model once

deployed requires an update after some time.¹ The reason for that is mainly Concept Drift [4], which may be summarized as a change of the response surface of the underlying problem over time [5]. Although Concept Drift is rather uncommon in pure academic tasks, it is of major importance in certain real-world applications. To emphasize that: in semiconductor fabrication, it is one main task for countless process and equipment experts to react on drifting processes or process machines.

A paper that garnered considerable attention was published by Gama et al. in 2014 [6], introducing a taxonomy for the subject area which is used in this article. This was followed by additional works summarizing various techniques for adapting to concept drift: In their 2019 review on Concept Drift Adaptation (CDA), Lu (and Gama) et al. analyzed 130 papers that included adaptation techniques [7]. They discovered that there are techniques which utilize retraining with varying window sizes, ensemble techniques used for adaptive purposes, and model-specific techniques like adaptive decision tree models. Another survey on Machine Learning for recurrent concept drift was conducted by Choudhary et. al. [8]. They listed additional methods that can be grouped into 1) The use of strategies for selecting, weighting or multiplying specific data points for retraining [9] [10] [11] and 2) Dynamics through adaptation in preprocessing [12], [13]. Suárez-Cetrulo et. al. additionally list Meta-learning techniques in the context of Concept Drift adaptation [5], p.9 ff. Meta-Learning has also been a topic that has continued to gain attention. Vanschoren [14] provides an overview of the Meta-Learning concept and its application areas. Among other Meta-Learning applications, they deal with Meta-models that “recommend the most useful [algorithm] configurations [...] given the Meta-features [...] of a task [...]” ([14], p.10). One option to do this is Transfer Learning ([14], p.12). With this, optimized models from previous tasks are taken as a warm-start to be trained on

*These authors contributed equally to this work and are listed in alphabetical order.

¹Model updates are part of the maintenance phase in context of the Cross Industry Standard for Data Mining (CRISP-DM [3]).

new tasks and yield a new model optimized for the new task.

The intersection between Meta-Learning and Concept Drift adaptation was reviewed in a survey published in 2023 by Son et. al. [15]. It contains several Meta-Learning techniques primarily used for the adaptation of Deep Learning models. In this development, it is clear to see that the topic of model adaptation has received a lot of attention and there have been continued efforts put into developing sophisticated techniques for model adaptation (for example [16], [17], [18], [19]).

In summary:

For a real-world deployment of a Machine Learning model that suffers from Concept Drift, there has to be a model update implementation to keep the application usable.

Consequently, there is a demand for techniques that can effectively update regression models. Furthermore, when multiple update techniques are available, it becomes important to determine which technique is most suitable for a specific use case and model type.

To qualitatively understand the term “most suitable”, two constraints are taken into account: 1) the required model accuracy, and 2) the availability of new labeled data.

The question then arises:

Which update technique extracts information from the new labeled data and integrates them into the model in the most efficient way?

While the inclusion of the accuracy requirement is obvious, the availability of new labeled data may need more clarification. Since new labels are expensive to obtain, they usually are not available in a high (enough) frequency. An example in the presented use case from Virtual Metrology are maintenances of the process machine: it is not uncommon that the whole distribution of a feature value jumps to another unknown(!) distribution mean due to the maintenance. So, the model’s predictions after machine maintenance will not be accurate enough. Since the frequency, of which new labeled data is available, is low, the data collection for a model update takes long. Hence, the aim is to adapt the model to fit the accuracy requirements with as few new data points as possible.

It should be stated clearly that this work does neither focus on the origin of the Concept Drift, nor what the Concept Drift looks like. Only the adaption of the model to fit the accuracy requirements is regarded.

There is one additional question that shall be addressed: Why are only model-agnostic update techniques regarded in this work?

The answer may be a bit uncomfortable: In the underlying industrial environment, the utilization of very specific techniques, e.g. a special type of Machine Learning Model and corresponding to that, a special model update technique, scales badly. It is much more beneficial to be flexible regarding the used Machine Learning model types. Especially, when the code bases of Machine Learning libraries get updated or completely new model types are developed, the deployed model update setup should still work.

Given the multitude of potential use cases and the variety of model types, one would want to apply update techniques

that can be scaled to different model types. It eliminates the need to start from scratch with each model, thereby making it economically more viable².

There are also other circumstances that lead to additional requirements on the techniques used: Typically, there is a scarcity of historical data for a given use case or process. This is primarily due to the high cost associated with retrieving labeled data, such as physical measurements, which are the labels for the supervised learning task in case of Virtual Metrology. Moreover, this work posits that the update techniques should also be applicable to newly introduced use cases or models which solve similar underlying problems. The idea is to let the model reuse already learned knowledge from the old case and relearn similar patterns again just based on the new case’s data.

This paper compares existing techniques to update regression models by applying them to a real-world use case from semiconductor fabrication (deposition of a layer on a wafer). They are also applied to an artificial use case to compare the behavior of the techniques in a setting where no unknown effects are present. As this work focuses on supervised regression tasks, the techniques considered here should be developed for regression tasks or be straightforwardly applicable to them. Extending update techniques from classification models to regression models would be a topic for “future work.” With the learnings in this work, a starting point for scaling update techniques across use cases with different model types in fabrication is provided.

The remainder of this article is structured as follows: section I-A provides an overview of the related work. The setup of the experiments is explained in section II. There it is also listed which update techniques are used. Section III presents the results with discussion. A summary and conclusion is given in section IV.

A. Related Work

A range of adaptive models can be found in literature. Examples include Adaptive Random Forest [20], Adaptive XGBoost for classification [21] and regression [22], and Adaptive Support Vector Machines [23] and others [24]. Neural Networks are inherently incremental learners, and there are works that introduce structures to Neural Networks to handle Concept Drift. For instance, Memory-Augmented Neural Networks [25] and Adaptive Extreme Learning Machines [26] for classification and [27], to name a few.

Why not use adaptive models? The choice of model type is specific to the use case. From experience in modeling processes in fabrication, Deep Learning models are often not utilized, primarily due to insufficient amounts of data to train them (costly data). In the context of this paper’s objective, which is to gain insights into model adaptation techniques that may be transferable to other use cases, it is sensible to investigate techniques applicable to any model type.

Update techniques that can be applied to any model type may be categorized, in the authors’ view, into 1) Selection

²Data Science resources are limited

and weighting of data points given to the model for updating (whatever mechanism the model then uses to update itself) [28], [29]. 2) External optimizers – that is, the internal model parameters are set to new values by an external optimizer and then reassigned to the model [30], [18], [19]. 3) Ensemble Methods [31].

The three cases are discussed further in the following:

1) Selection and weighting of data points for updating: Rolling Window [6] involves different window sizes and weighting of the input data in the loss function. This can be done linearly or exponentially decreasing against the age of the data. There are also newer techniques for determining the weights of the data points, or even leaving out data points entirely. For instance, [32] uses fuzzy models to select which data are used. Similarly, [33] employs fuzzy kernel c-means clustering of the data stream to determine the data points for the update. This is combined with a forgetting policy during the update. The techniques included in this paper should be applicable with as little lead time and thus as little historical data as possible. Given that it is assumed that there will generally not be enough historical data available for the auxiliary models for selecting update data for the regression model, these techniques are not applied here. The same applies to techniques that store a “model history” and make a selection with various strategies as to which model or which composition of historical models should be used for new inference data.

2) External optimizers: What can an external optimizer do better than a model’s own one? Possibly, information from past data can be stored in an external optimizer. This happens in Meta-Learning. An example of this is Andrychowicz’ technique “Learning to Learn by Gradient Descent by Gradient Descent”³ to optimize an algorithm using a Meta Learner [30]. As optimizer (in this case called Meta-Learner), they use Long Short-Term Memory (LSTM) networks [34]. The LSTM’s task is to predict the changes in the algorithm’s parameters given the gradient of the loss of the task of the algorithm with respect to the algorithm’s parameters. This means, the LSTM predicts the changes in the Neural Network’s internal parameters (weights and biases) given the regression/classification loss’ gradient with respect to the weights and biases of the Neural Network. In their paper, the training of Neural Networks on the MNIST [35] and CIFAR-10 [36] dataset are used to demonstrate the developed architecture. Li and Malik [19] developed a similar approach, formulating the problem in a reinforcement learning setting.

Another Meta-Learning approach is Model-Agnostic Meta-Learning (MAML) [37]. Here, starting values are sought for the internal model parameters, from which the best accuracy is achieved when using a fixed optimizer to do the adaptation to new data. The key idea behind MAML is to find a model initialization that is not only good for one task but can be quickly fine-tuned for any task in the distribution of tasks. During the training phase, MAML performs a two-level

optimization process. In the inner loop, it learns a separate model for each task using the inner optimizer. For example, the optimizer can use several gradient steps. In the outer loop, MAML updates the initial model parameters based on how well the task-specific models performed. The goal is to find a set of initial parameters that, when fine-tuned on a few examples from a new task, can achieve good performance on that task. Finn et. al. [37] demonstrate their technique on Neural Networks solving an artificial regression use case (sine wave) and classification tasks.

The Meta-Learning techniques mentioned above were applied in the Deep Learning field [15]. The techniques use gradients of the loss function wrt. internal model parameters in the optimization processes [30], [37], [18]. The name of the technique “Learning to Learn without Gradient Descent by Gradient Descent” [38] may suggest that computing the gradients of the loss function with respect to the internal parameters of the model to be optimized is omitted. Indeed, the gradients are not fed into the Meta Learner during inference phase. But to train the Meta Learner, still, the gradients mentioned are necessary.

For Neural Networks or much simpler multilinear regressors, for example, these techniques are straightforwardly applicable. Theoretically, the application of the techniques or ideas could also be extended to models with discontinuities in the loss function by finding ways to artificially make the loss function continuous (e.g., overlaying of a smoothing function) or other numerical calculation of the gradients. Extending the techniques that work well in the Deep Learning field to other model types such as Trees / Forests / GXBoost is an interesting task. However, this is a separate work. As an example, two of these techniques are nevertheless applied to the Neural Network in this article.

3) Ensemble techniques. In the context of ensemble techniques, there exists the utilization of multiple model types and the (weighted) aggregation of various results [39],[40],[41],[42]. The construction of different models based on diverse input data and the subsequent aggregation of results is referred to as Bagging, with Adaptive Bagging also being a notable variant [31]. An advanced approach for the latter two involves predicting the weights with a meta-model. However, this approach is not adopted here due to the general lack of sufficient training data at the onset. Another ensemble technique is Boosting, where the next model is fitted on the residuals.

This paper applies a selection of model update techniques to a real-world regression task from a use case in semiconductor fabrication and compares these update techniques, which are practically applicable. The selection of model update techniques is based on the industry requirements discussed in Section I. Furthermore, it is a requirement that the techniques can be applied as best as possible to newly introduced regression models.

The literature provides two studies on the application and comparison of update techniques:

Celik et. al. investigate update strategies for Auto-ML

³This update technique is referred to as “grad2-lstm” in this article.

systems in the face of Concept Drift [43]. The strategies employed encompass various configurations of “Detect and Restart” vs. “Periodic Restart” and “warm-start” vs. “re-train” vs. “AutoML-Restart”. The difference between “retrain” and “AutoML-Restart” is that “AutoML-Restart” includes re-tuning of hyperparameters. A not unexpected result, quote [43]: “different drift characteristics affect learning algorithms in different ways, and that different adaptation strategies may be needed to optimally deal with them”. This indicates that there is no one-for-all technique regarding model adaptation.

The work at hand does not focus on Auto-ML systems but looks at a level below, to see the direct effect of different update strategies.

In their benchmark on “Learning to Optimize” (LTO), Chen et. al. [44] compare optimization techniques on various benchmark tasks. One of these tasks is the training of Neural Networks. The concrete optimization tasks are training a Multilayer Perceptron and a Convolutional Neural Network with different optimizing techniques on MNIST data set [35]. For both experiments, as traditional optimizers, Adaptive Moment Estimation (ADAM) [45], SGD [46] and Root Mean Square Propagation [47] are used. Four Learning-to-optimize approaches are used in the experiments [30], [48], [49], [50]. For all of them, the architecture of the optimizer contains recurrent Neural Networks such as LSTM architectures. Among these is the approach from Andrychowicz et. al. [30], which is used in the work at hand as well. The findings of Chen et. al. for the experiments of training Neural Networks are “lacking stability during testing” ([44], p.35) for all of the LTO-optimizers whereas the traditional optimizers show convergence when training each Neural Network. Training the Multilayer Perceptron, the optimizer from [30] diverges. It works for training the Convolutional Neural Network, but only for a small number of iterations. Two of the other three LTO-approaches show good results on the Multilayer Perceptron but can not do efficient optimization on the Convolutional Neural Network [48], [49]. The use of an advanced training scheme for LTO-Optimizers [50] allowed [44] to improve the optimizing performance on both Neural Network optimizers.

It is worthy to mention that the two experiments from [44] were carried out on Neural Network optimizers that performed classification of the MNIST [35] dataset. The article at hand aims at ongoing optimization of models for regression tasks (which can be but are not necessarily Neural Networks). For our experiment with a Neural Network performing regression tasks, [30] is used as a first benchmark technique on the given regression data.

To the best knowledge of the authors, there is no comparison work yet that focuses on model type independent update techniques and compares their functionality on different model types.

Especially, a comparison of different update techniques on different regression models (model types) on the same real-world regression task is important but not available. This paper makes a starting point for closing this gap. The goal is to gain

knowledge about which update techniques perform better on which model type in an industrial regression setting.

II. SETUP OF THE COMPARISON

To evaluate the performance of the different model update techniques, multiple experiments are conducted, where a model is trained on a given data set and afterwards used with drifting test data. During the test, the model is continually updated with latest test data using the different update techniques. Two independent cases are regarded: the first one is a real-world application from an industrial environment and the second one is based on artificial data to get an undisturbed view on the update techniques.

A. Real-World Use Case

The industrial use case is Virtual Metrology of a layer deposition in semiconductor fabrication. The regression data set consists of 32 features (numerical and encoded categorical) and one label. A feature selection process was done before this. The label (and prediction) data are the thicknesses of a layer deposited on a wafer and are scaled to maintain confidentiality.

For the Virtual Metrology use case, the available labeled data was divided into 165,000 data points for training and 250,000 data points for the comparison of the update techniques. The data points for the comparison were divided into batches of 50 data points yielding 5000 batches. Two models were trained including hyper parameter optimization: An XGBoost model and an Artificial Neural Network⁴ ([51], [52]). After the training, batches of feature data points are fed to the models in historically correct order and compared to the corresponding label data. One batch contains 50 wafers (data points). Each batch is handled as follows:

A train-test-split is made yielding $0.6 * 50 = 30$ wafers for updating and $0.4 * 50 = 20$ wafers for evaluation of the updated model. The latter provides the accuracy for the given batch. As accuracy measure, the mean absolute error (MAE) is taken. For each update technique, for each batch, the accuracy is stored. To make a statistically valid statement about the accuracy, over the stored accuracies, the mean of the accuracies in the window is calculated, using a rolling window of size 500.

The question asked is which update techniques work how well on different model types. This is assessed by using two setups: blind adaptation and informed adaptation [6]. Blind adaptation means the model is updated regularly without a trigger, whereas informed adaptation means the model is updated when triggered, e.g. via a Concept Drift detector or a performance degradation detector. For the blind adaptation setting, the model is updated with the corresponding update technique regularly in intervals of 200 batches. The frequency of updating and the update batch size were not varied because it is not the focus of this work. As loss function the mean absolute error is used.

⁴Sequential, dense Neural Network with the structure: Input Layer of size 32, Dense Layers with [64,128,64,64,32,32,1] nodes. Each layer is followed by a ReLU activation function.

For the informed updating case, in this work, a performance degradation detector is used: the model accuracy (here: the model loss on available new labeled data) exceeding a fixed threshold. In practice, this threshold is use-case specific. In this study, on the same real world use case, the experiments are conducted with two different model types: XGBoost and Neural Network. The accuracy of the initially trained XGBoost model is higher than the initial accuracy of the Neural Network. In practice, for this use case, the XGBoost model would be chosen. For academic purposes, possibly allowing to gain knowledge for other use cases, in which a Neural Network might achieve better performance, the experiments are conducted on both model types. The threshold used for the performance degradation detector is oriented at the initial accuracy of the corresponding model for the use case, leading to a higher threshold for the Neural Network than for the XGBoost model. It is important to mention that in practice, there would only be one threshold per use case. If a model would exceed this threshold in the beginning, it would not be, and is not, chosen.

To maintain confidentiality, neither the absolute values of the model performances nor the threshold values are given. The loss values, shown in the III section, are scaled to relative units.

Quantification, which update technique performed best, was done by calculating the overall mean loss over the whole 5000 batches.

The results are shown and discussed in Section III.

B. Generation of the Artificial Example Data

To validate the implementation of the adaptation techniques without any unknown influences, artificial training data were constructed. They consist of five numerical input data where each input lies in the interval $[0, 1]$. Independently, uniformly randomly sampling each variable from this interval, one corresponding numerical output (label) is generated using a multilinear function with known, fixed parameters a_1, a_2, a_3, a_4, a_5 and b :

$$f : [0, 1]^5 \longrightarrow \mathbb{R} \quad (1)$$

$$(x_1, x_2, x_3, x_4, x_5) \rightarrow f(x_1, x_2, x_3, x_4, x_5) \quad (2)$$

$$f(x_1, x_2, x_3, x_4, x_5) := \sum_{i=1}^5 a_i \cdot x_i + b \quad (3)$$

$$a_i \in \mathbb{R} \text{ constant for } i \in \{1, 2, 3, 4, 5\} \quad (4)$$

For the concrete example used here, $a_1 = 1.0$, $a_2 = 0.5$, $a_3 = -0.1$, $a_4 = -1.5$, $a_5 = 0.005$ and $b = 0.7$ were chosen for the pretraining data generation. The pretraining data consist of 200 batches (10,000 data points). For the comparison of the update techniques during inference, a data set – similar to the training data – is generated, which includes Concept Drift. The inference data consist of 20 batches containing 50 data points each. The inference data generation is done by changing the parameters a_3 and b three times, every fifth, batch, in the following way: The first five batches of the inference data were generated exactly from the same function as the pretraining

data. The data in the sixth batch are generated in the same way but changing the parameters in function (4) to

$$a'_3 = a_3 + \Delta a_3, \quad \Delta a_3 = 0.2, \quad b' = b + \Delta b, \quad \Delta b = 0.2.$$

Five batches of the new configuration were appended to the inference data. Then again, starting from this parameter configuration data generating function (equation 4), the function parameters are shifted again by $\Delta a_3 = 0.2$ and $\Delta b = 0.2$. This procedure is repeated one more time so that finally there are three sudden slight drifts in the inference data which are separated by five batches.

A Dense Artificial Neural Network⁵ was used to fit the artificial data because for Neural Networks, the update techniques *grad2-lstm* and *MAML* can also be applied.

The results are shown and discussed in Section III.

C. Model Adaptation Techniques used

The following model adaptation techniques are used in the comparison experiments.

- *baseline*. Refers to the original pretrained model with no update taken.
- *only-latest*. Using the model's retraining mechanism for the update and using only the data from the latest incoming data batch for the update.
- *rolling-window*. Using the model's retraining mechanism for the update and using data from the latest batch plus the data from the previous batches with a fixed sliding window size [28].
- *exp-forgetting*. Using the model's retraining mechanism for the update and using data from the latest batch plus the data from the previous batches with a fixed sliding window size. The importance of the data points decreases exponentially the older the data points are [29]. The decay factor is a hyperparameter that was optimized for the real and artificial use case separately.
- *boosting*. Boosting sequentially with the following setup: Take the original model and append a "booster model." Update only the booster model, Using a sliding window. For the real use case, in case of a Neural Network as base model, a small XGBoost model was chosen as the booster model. In case of an XGBoost model as the base model, a small XGBoost model was chosen as the booster model was chosen as well. For the artificial use case, a linear model was chosen as the booster model. The choices of the booster model types were made based on which booster model candidate yielded the better performance: Linear Regressor or XGBoost.
- *grad2-lstm*. The update technique "Learning to learn by gradient descent by gradient descent" [30]. This technique was only used with Neural Networks due to the limitations discussed section I-A (use of the gradient of

⁵Sequential, dense Neural Network with the structure: Input Layer of size 5, Dense Layer with 2 nodes and ReLu activation function, Dense Layer with 2 nodes and ReLu activation function and output layer with 1 node and ReLu activation function.

TABLE I
HYPERPARAMETERS FOR EACH UPDATE TECHNIQUE USED.

Update technique:	Update method's hyperparameters using XGBoost	Update technique's hyperparameters using NN
only-latest-WS	learning rate	learning rate, epoch number
only-latest-SC	learning rate	-
rolling-window-WS	learning rate	learning rate, epoch number
rolling-window-SC	learning rate	learning rate, epoch number
exp-forgetting-WS	learning rate, decay factor	learning rate, epoch number, decay factor
exp-forgetting-SC	learning rate, decay factor	learning rate, epoch number, decay factor
boosting	-	learning rate when XGB as booster
grad2-lstm	-	learning rates inner optimizer and outer optimizer, epoch number
MAML	-	number of gradient steps, learning rate outer optimizer, learning rate for gradient steps
baseline	-	-

the regression model's loss function with respect to the regression model's internal parameters).

- *MAML*. This update technique uses "Model-agnostic Meta learning" [37], using Adam optimizer for outer optimization and gradient descent for inner optimization. This technique was only used with Neural Networks due to the limitations discussed in section I-A (use of the gradient of the regression model's loss function with respect to the regression model's internal parameters).

For the model's retraining mechanism in case of Neural Networks, ADAM optimizer is used [53]. For the update techniques only-latest, rolling-window and exp-forgetting, the sliding window size was chosen same as the size of the pretraining data (165,000 data points). For MAML and grad2-lstm, a sliding window of 5,000 data points was used as increasing the number of update points did not lead to better accuracy. Two variants, warm-start and training from scratch, were conducted for the update techniques only-latest, rolling-window and exp-forgetting. One exception is only-latest in case of the Neural Network since it is theoretically trivial that training a Neural Network of the given size on only 30 data points will not work.

The hyperparameters of the update techniques were optimized per use case and per model. The existing hyperparameters for each technique are shown in Table I. The optimization of hyperparameters was conducted within the batch range of 0 to 3,000 batches, as opposed to the full range up to 5,000. This approach was adopted based on the assumption that continual hyperparameter optimization for update techniques will not be feasible in subsequent real-world applications. In the tables and figures of this work, the terms "warm-start" and "training from scratch" are abbreviated by "WS" and "SC" respectively.

III. RESULTS

The findings that were made in the course of applying the chosen model update techniques to the industrial use case are presented in the following.

TABLE II
OVERALL LOSSES FOR ALL EXPERIMENTS CONDUCTED.

Use case Regression model Update mode	Artificial case NN		Virtual Metrology NN		XGB	
	Infm	Blind	Infm	Blind	Infm	Blind
Update technique:						
only-latest-WS	0.06	0.06	0.72	0.71	2.17	3.73
only-latest-SC	-	-	-	-	2.73	3.64
rolling-window-WS	0.46	0.46	0.13	0.06	0.41	0.51
rolling-window-SC	0.49	0.47	0.39	0.33	0.34	0.40
exp-forgetting-WS	0.46	0.46	0.22	0.09	0.43	0.49
exp-forgetting-SC	0.46	0.47	0.43	0.39	0.32	0.39
boosting	0.12	0.19	0.69	0.68	0.92	0.92
grad2-lstm	0.15	0.17	0.76	0.76	-	-
MAML	0.10	0.18	0.71	0.89	-	-
baseline	0.47	0.47	0.72	0.72	0.63	0.63

TABLE III
THE VALUES SHOW HOW OFTEN THE CORRESPONDING UPDATE TECHNIQUE WAS UPDATED IN THE GIVEN USE CASE FOR THE INFORMED ADAPTATION SETTING.

Update technique:	Artificial UC	Real UC NN	Real UC XGB
only-latest-WS	3	24	440
only-latest-SC	-	-	739
rolling-window-WS	13	20	89
rolling-window-SC	13	15	83
exp-forgetting-WS	13	17	81
exp-forgetting-SC	13	17	83
boosting	10	20	175
grad2-lstm	13	23	-
MAML	6	26	-
baseline	0	0	0

Table II shows the results for the experiments conducted.⁶ Two use cases, artificial and real (Virtual Metrology), were regarded. For both use cases, Neural Networks were used and additionally, for the real-world use case an XGBoost model was used. The loss (mean absolute error) for each subsequent batch was scaled to a reference in order to maintain confidentiality. This is defined in the following way: The normalized loss equals zero for the lowest loss of the baseline (no model updating, black curves). The normalized loss equals one for the maximum loss of the baseline. This means, the losses are shown in relative units (relative to the losses of the baseline, per use case and per model). The mean of the losses in relative units, taken over the whole inference period of 5,000 batches, is referred to "overall loss" in this article.

In Table III it is shown for the informed case, for which model update techniques how many updates were performed during the inference (per use case and model).⁷ The blind adaptation with fixed update frequencies (updating every 5 batches for the artificial use case and every 200 batches for the real use case) led to 3 updates for the artificial and 24 updates for the real use case, for each update technique. Comparability of the values in tables II and III is given column-wise, not row-wise, due to the experiment setup.

⁶Informed is abbreviated by "Infm."

⁷Use Case is abbreviated by UC. XGBoost is abbreviated by XGB.

A. Neural Network Model for Artificial Use Case

For showing and validating the implemented update techniques on data with no unknown influences, artificial data were generated. Three times, every fifth batch, a Concept Drift was incorporated. Details about the data construction are explained in Section II-B.

1) *Informed adaptation*: The evolution of the loss of the Neural Network during inference for the informed update setting is shown in Fig. 1 for all update techniques applied. Model degradation, being caused by the Concept Drift, can clearly be seen in the baseline curve's increasing loss (black curve). The rolling window techniques, whether training from scratch or using warm-start, perform similar or worse than not updating at all. This is because the Concept Drifts are introduced suddenly and there is no recurrence in it by construction of the data. In the data set used for updating, the data from the newest batches are too underrepresented to make an improvement in accuracy.

Exponential forgetting warm-start (red curve) brings a slight improvement for each successive batch after a drift was introduced.

An improvement and thus successful handling of the Concept Drift is achieved by the update techniques "only-latest-WS," "grad2-lstm," "MAML" and "boosting". Using warm-start and tuning to only the latest batch, "only-latest-WS" can adapt best to the new data. For all four update-techniques, at batch numbers 5, 10 and 15, a peak is observed in the loss curves. This is due to the experiment setup: As detector of drift, exceeding of an accuracy threshold is used. When this is detected for a batch, the update and evaluation of the updated model is performed on the next batch.

The "boosting"-technique also uses only the latest batch for updating. The booster model (model in the sequential ensemble that is closest to the output) only is updated here.

Model-Agnostic Meta Learning (MAML) gave second best results. It took two subsequent batches to update, until MAML achieved the same accuracy as "only-latest-WS." "grad2-lstm" did not achieve this accuracy, but still maintained loss values that showed clear improvement wrt. the baseline.

At this point, the correctness of the implementation especially of the techniques "MAML," "grad2-lstm" and "boosting" is verified.

2) *Blind adaptation*: For the case of blind adaptation, the same setting as in the last subsection is used, only that there is no trigger for the update. The numbers of the batches at which an update is performed are predefined and set to 6, 11 and 16 (for each update technique). With this setting, updating takes place one batch after the drift occurs. The resulting difference to the informed case is that after a drift occurred only one update is performed for each technique, until the next drift occurs.

The results are shown in Fig. 2. It is visible that for "MAML," "grad2-lstm" and "boosting," this leads to higher losses as compared to the informed case. This means these update techniques do profit from carrying out the additional




Fig. 1. Informed adaptation for the artificial use case using a Neural Network. WS= warm-start, SC= from scratch. The overall losses for each update technique are given in table II, column 2.




Fig. 2. Blind adaptation for the artificial use case using a Neural Network. The overall losses for each update technique are given in table II, column 3.

updates as in the informed case. Still, for these techniques, one update suffices to see a clear improvement as compared to the baseline of not updating at all.

B. Virtual Metrology using a Neural Network

Having validated the implementation of the various update techniques and observed their behaviour on an artificial use case with known data properties and drifts, a real use case with unknown drift characteristics is regarded. The same update techniques as for the artificial use case are applied. Each update technique's hyperparameters were tuned for this use case.

1) *Informed adaptation*: The loss curves for the informed updating experiment for the VM case, using a Neural Network as regression model, are shown in Fig. 3. First of all, the model

degradation with time can be seen in the black curve (baseline, not updating) for which the loss increases during the inference experiment with incoming batches.

As opposed to the artificial use case, in which rolling window and exponential forgetting performed worst, in the real world use case they perform best. Especially the "rolling-window-warm-start" technique, which is the best-performing technique in this setting, is capable of restoring the initial loss.

The update technique "only-latest" which, using warm-start, fine tunes the model on the latest data batch only, results in nearly the same losses as the baseline. The "boosting" method, which also uses only the latest batch for updating, performs slightly, but not much better than the baseline (overall loss of 0.69 compared to baseline 0.72).

This shows that it is crucial for this use case to keep data points from former batches in the set of data points for updating (3300 data batches were used as sliding window). The number of batches for updating (updating-batch size) was not varied further since this is not the focus of this work.

Both techniques Model-Agnostic Meta Learning and Learning to learn by gradient descent by gradient descent, did not bring an improvement. For both cases, the optimization of the hyperparameters yielded hyperparameters that correspond to nearly not updating the model. The overall loss of grad2-lstm is slightly worse than not updating for both settings blind and informed. The loss of the MAML method is comparable to the not updating baseline for the informed case and slightly higher than the baseline for the blind case. For both, MAML and grad2-lstm, there are ranges in which they perform slightly better than the baseline. But these ranges are the exception from the overall range. The conclusion here is that for this use case, for the used setting of update frequency and batch size, the two methods do not bring improvements. In the grad2-lstm case, increasing the size of the LSTM also did not lead to any improvements.

2) *Blind adaptation*: In the blind adaptation setting, the rolling window warm-start technique performs best, being followed by the exponential forgetting warm-start technique (cf. Fig. 4). The difference between these two methods' overall losses is not as high as in the informed case.

The differences are likely to originate from the different update frequencies used in the blind adaptation case than those in the informed setting, resulting from excess of the accuracy threshold. Lowering the accuracy threshold for the Neural Network informed case, would presumably result in lower overall loss values. The dependency of the performance of the different update techniques on the update frequency therefore seems to be a question worth asking. For follow-up work, together with varying the update-batch size, this would be an interesting study.

The remaining update techniques qualitatively have the same relative overall losses to the exponential forgetting warm-start and rolling window warm-start techniques as well as to each other. The same conclusions that were drawn in the last subsection (informed case), hold.




Fig. 3. Virtual Metrology regression with Neural Network – Informed adaptation setting. The overall losses for each update technique are given in table II, column 4.

To make the different lines more distinguishable, every 15th data point was plotted large. However, the size of the data points has no meaning.




Fig. 4. Virtual Metrology regression with Neural Network – Blind adaptation setting. The overall losses for each update technique are given in table II, column 5.

To make the different lines more distinguishable, every 15th data point was plotted large. However, the size of the data points has no meaning.

C. XGBoost model for Virtual Metrology

1) *Informed adaptation*: For the XGBoost model in the informed case (Fig. 5), the method exponential forgetting from scratch gave the best overall loss. It was followed by the method rolling window from scratch. The corresponding techniques with warm-start performed slightly worse. The "boosting" technique did not lead to improvement on the model performance with respect to the baseline but even increased the losses. Using only the data from the new batch for the model retraining ("only-latest"-techniques), made the performance even worse. This suggests that the reason for the




Fig. 5. Virtual Metrology regression with XGBoost model – Informed adaptation setting. The overall losses for each update technique are given in table II, column 6. To make the different lines more distinguishable, every 30th data point was plotted large. However, the size of the data points has no meaning.




Fig. 6. Virtual Metrology regression with XGBoost model – Blind adaptation setting. The overall losses for each update technique are given in table II, column 7. To make the different lines more distinguishable, every 30th data point was plotted large. However, the size of the data points has no meaning.

low performance of "boosting" in the real use case is related to the data and not to the update strategy "boosting" in general. One batch seems not to represent the data sufficiently.

2) *Blind adaptation*: Similarly to the informed adaptation case, the best-performing update technique is exponential forgetting starting from scratch, being followed by rolling window starting from scratch and the two corresponding methods in warm-start setting (cf. Fig. 6). The boosting approach has the same overall accuracy as the not-updated model. Using only the latest batch for updating is not recommended as well.

IV. SUMMARY AND CONCLUSION

This work compares model-type-independent update techniques on regression models. For this, a regression use case from semiconductor fabrication (Virtual Metrology (VM) of a layer deposition process), was used. To validate the functionality of the update methods, a simple artificial use case (data sampled from a multilinear function) was used. On the VM use case, experiments were done with an XGBoost model and with a Neural Network. For all of these cases, a blind and an informed updating strategy was applied, yielding six scenarios that are shown in table II. Eight update techniques were regarded. All of them were successfully validated on the artificial use case. Using the Neural Network’s own warm-start fine tuning method for updating, using only the latest batch, performed best in the artificial case in both, the informed and blind setting.

For the real use case (VM), using only the newest batch for updating performed poorly in all scenarios. For both, the XGBoost and Sequential Neural Network, sliding window techniques performed best in the VM case. In case of the XGBoost model, it did not make a significant difference if training from scratch or warm-start was used. For the Neural Network, warm-start clearly performed better than training from scratch. This is explainable by the good transfer learning capability of Neural Networks ([14], p.12). The boosting approaches lead to slight improvement for the Neural Network and to no improvement for the XGBoost model.

The best technique for VM, XGBoost, was exponential forgetting, training from scratch. For the Neural Network in the VM case, rolling window with warm-start performed best.

In this work, the batch size used for updating and the updating frequency were set to constant values. Nevertheless, comparing the performances of informed and blind scenarios for the real use case leads to the (obvious) conclusion that which update method performs best, depends on the frequency of the updates as well as the batch size of the updates. A study varying these hyperparameters would be an interesting follow-up work.

For the Neural Network, two Meta-learning update techniques were used in this work: "Model-agnostic meta learning" and "Learning to learn by Gradient descent" [37], [30]. Both of them were not able to improve the loss as compared to the base line. It is important to note that their implementation was validated on the artificial use case, on which they showed significant improvement. Chen et. al. conducted an experiment in which they trained a) a Multilayer Perceptron and b) a Convolutional Neural Network on the MNIST dataset, using different optimizers [44], [35]. Among these optimizers, the approach "Learning to Learn by Gradient Descent" from Andrychowicz, called grad2-lstm in this article, was used as well. In the experiment of [44], on the Multilayer Perceptron, the optimizer did not converge. This result is similar to the result that was obtained in the article at hand for the grad2-lstm method, where it did not provide loss improvement in the real world use case. The experiment from [44] worked for

the Convolutional Neural Network they used, but only for a small number of iterations. It is important to emphasize that, in the work at hand, no Convolutional Network is used and the task that the optimizee performs is not classification, but regression.

Although the Meta-Learning techniques used, MAML and grad2-lstm, did not show improvement for this real use case and setting (meaning fixed updating batch size and updating frequency), it might be good to not discard them when doing studies with varying batch size for updating. This is because they might outperform other techniques for smaller updating batch sizes (in the case of few-shot learning).

The insight gained from the experiments conducted is there is no "one update technique" which performs best in all experiments. For the Virtual Metrology use case, the group of sliding window techniques has superior performance. For the artificial use case, this group performed worse than other techniques. For the artificial use case, the bad performance of the sliding window techniques is explainable by the drift properties: The Concept Drift incorporated into the artificial use case consists of slight sudden jumps in the function generating the training data. Such a behaviour can be theoretically expected in real use cases as well, due to maintenance events. In practice, not only maintenance events are the cause of drifts, but a variety of other influences might change the response surface as well over time. To explain why sliding window techniques performed well on the real use case, and others performed worse (for instance MAML and grad2-lstm), the properties of the real use case's data and their drift properties could be beneficial. The examination of the use case's data and drift properties were not the focus of this work though. This work aims at finding the most suitable update technique for the given use case for different possible models. This is because scalability to other use cases, for which different model types might be best-suited, is desirable. The focus therefore lied on examining the performance (accuracy) of model-independent update techniques. More knowledge about scalability of update methods to other use cases and models can be gained by performing comparable studies on other real world use cases, for example other Virtual Metrology cases. It would be possible to examine, if for each use case the same update techniques perform best, and if not, to ask what causes the differences. Presumably, it might be helpful to examine data and drift properties of the different use cases in course of this, which makes this topic an interesting task for future work. With a wider range of experiments (on more use cases, with more model types and with variation of update batch size and update frequency), recommendations for new use cases (potentially using different model types) could be given and thus scaling of the update technique application can be aided. The work at hand provides a starting point for this research.

ACKNOWLEDGMENT

This work is funded by the European Union within "NextGeneration EU", by the Federal Ministry for Economic Affairs and Climate Action (BMWK) on the basis of a decision

by the German Bundestag and by the State of Saxony with tax revenues based on the budget approved by the members of the Saxon State Parliament in the framework of "Important Project of Common European Interest - Microelectronics and Communication Technologies", under the project name "EUROFOUNDRY".

REFERENCES

- [1] V. Maitra, Y. Su, and J. Shi, "Virtual metrology in semiconductor manufacturing: Current status and future prospects," *Expert Systems with Applications*, vol. 249, p. 123559, 2024. doi: 10.1016/j.eswa.2024.123559
- [2] S. Yan, C. Luo, S. Wang, S. Ding, L. Li, J. Ai, Q. Sheng, Q. Xia, Z. Li, Q. Chen, S. Li, H. Dai, and Y. Zhong, "Virtual metrology modeling for cvd film thickness with lasso-gaussian process regression," in *2023 China Semiconductor Technology International Conference (CSTIC)*, 2023. doi: 10.1109/CSTIC58779.2023.10219236 pp. 1–4.
- [3] C. Schröder, F. Kruse, and J. M. Gómez, "A systematic literature review on applying crisp-dm process model," *Procedia Computer Science*, vol. 181, pp. 526–534, 2021. doi: 10.1016/j.procs.2021.01.199
- [4] F. Bayram, B. S. Ahmed, and A. Kassler, "From concept drift to model degradation: An overview on performance-aware drift detectors," *Knowledge-Based Systems*, vol. 245, p. 1, 2022. doi: 10.1016/j.knsys.2022.108632
- [5] A. L. Suárez-Cetrulo, D. Quintana, and A. Cervantes, "A survey on machine learning for recurring concept drifting data streams," *Expert Systems with Applications*, vol. 213, p. 118934, 2023. doi: 10.1016/j.eswa.2022.118934. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957417422019522>
- [6] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–37, Mar. 2014. doi: 10.1145/2523813
- [7] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 12, pp. 2346–2363, 2019. doi: 10.1109/TKDE.2018.2876857
- [8] A. Choudhary, P. Jha, A. Tiwari, and N. Bharill, "A brief survey on concept drifted data stream regression," in *Soft Computing for Problem Solving*, A. Tiwari, K. Ahuja, A. Yadav, J. C. Bansal, K. Deep, and A. K. Nagar, Eds. Singapore: Springer Singapore, 2021. doi: 10.1007/978-981-16-2712-5_57 pp. 733–744.
- [9] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, "Generative adversarial networks: An overview," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 53–65, 2018. doi: 10.1109/MSP.2017.2765202
- [10] Y. Song, G. Zhang, J. Lu, and H. Lu, "A fuzzy kernel c-means clustering model for handling concept drift in regression," in *2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2017. doi: 10.1109/FUZZ-IEEE.2017.8015515 pp. 1–6.
- [11] D. Liu, Y. Wu, and H. Jiang, "Fp-elm: An online sequential learning algorithm for dealing with concept drift," *Neurocomputing*, vol. 207, pp. 322–334, 2016. doi: 10.1016/j.neucom.2016.04.043
- [12] S. J. Delany, P. Cunningham, A. Tsymbal, and L. Coyle, "A case-based technique for tracking concept drift in spam filtering," in *Applications and Innovations in Intelligent Systems XII*, A. Macintosh, R. Ellis, and T. Allen, Eds. London: Springer London, 2005. doi: 10.1007/1-84628-103-2_1. ISBN 978-1-84628-103-7 pp. 3–16.
- [13] Y. Song, G. Zhang, H. Lu, and J. Lu, "A noise-tolerant fuzzy c-means based drift adaptation method for data stream regression," in *2019 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2019. doi: 10.1109/FUZZ-IEEE.2019.8859005 pp. 1–6.
- [14] J. Vanschoren, "Meta-learning: A survey," *arXiv preprint arXiv:1810.03548*, 2018. doi: 10.48550/arXiv.1810.03548
- [15] J. Son, S. Lee, and G. Kim, "When meta-learning meets online and continual learning: A survey," 2023. doi: 10.48550/arXiv.2311.05241
- [16] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," 2019.
- [17] S. Lee, H. Jeon, J. Son, and G. Kim, "Sequential bayesian continual learning with meta-learned neural networks," 2024. [Online]. Available: <https://openreview.net/forum?id=6r0B0lb771>
- [18] J. von Oswald, C. Henning, B. F. Grewe, and J. Sacramento, "Continual learning with hypernetworks," 2022. doi: 10.48550/arXiv.1906.00695

- [19] K. Li and J. Malik, "Learning to optimize," 2016. doi: 10.48550/arXiv.1606.01885
- [20] H. M. Gomes, J. P. Barddal, L. E. B. Ferreira, and A. Bifet, "Adaptive random forests for data stream regression." in *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018. [Online]. Available: https://www.ppgia.pucpr.br/~jean.barddal/assets/pdf/arf_regression.pdf
- [21] J. Montiel, R. Mitchell, E. Frank, B. Pfahringer, T. Abdesslem, and A. Bifet, "Adaptive xgboost for evolving data streams," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020. doi: 10.1109/IJCNN48605.2020.9207555 pp. 1–8.
- [22] F. M. de Souza, J. Grando, and F. Baldo, "Adaptive fast xgboost for regression," in *Intelligent Systems*, J. C. Xavier-Junior and R. A. Rios, Eds. Cham: Springer International Publishing, 2022. doi: 10.1007/978-3-031-21686-2_7. ISBN 978-3-031-21686-2 pp. 92–106.
- [23] J. Zheng, F. Shen, H. Fan, and J. Zhao, "An online incremental learning support vector machine for large-scale data," *Neural Computing and Applications*, vol. 22, pp. 1023–1035, 2013. doi: 10.1007/s00521-011-0793-1
- [24] Łukasz Korycki and B. Krawczyk, "Adaptive deep forest for online learning from drifting data streams," 2020. doi: 10.48550/arXiv.2010.07340
- [25] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "One-shot learning with memory-augmented neural networks," 2016. doi: 10.48550/arXiv.1605.06065
- [26] S. Xu and J. Wang, "Dynamic extreme learning machine for data stream classification," *Neurocomputing*, vol. 238, pp. 433–449, 2017. doi: 10.1016/j.neucom.2016.12.078
- [27] N. Mishra, M. Rohanijad, X. Chen, and P. Abbeel, "A simple neural attentive meta-learner," 2018. doi: 10.48550/arXiv.1707.03141
- [28] B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom, "Models and issues in data stream systems," in *Proceedings of the Twenty-First ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, ser. PODS '02. New York, NY, USA: Association for Computing Machinery, 2002. doi: 10.1145/543613.543615. ISBN 1581135076 p. 1–16.
- [29] R. Klınkenberg, "Learning drifting concepts: Example selection vs. example weighting," *Intell. Data Anal.*, vol. 8, pp. 281–300, 2004. doi: 10.3233/IDA-2004-8305
- [30] M. Andrychowicz, M. Denil, S. Gómez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. de Freitas, "Learning to learn by gradient descent by gradient descent," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2016/file/fb87582825f9d28a8d42c5e5e5e8b23d-Paper.pdf
- [31] N. Oza, "Online bagging and boosting," in *2005 IEEE International Conference on Systems, Man and Cybernetics*, vol. 3, 2005. doi: 10.1109/ICSMC.2005.1571498 pp. 2340–2345 Vol. 3.
- [32] E. Lughofer, "Efficient sample selection in data stream regression employing evolving generalized fuzzy models," in *2015 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2015. doi: 10.1109/FUZZ-IEEE.2015.7337844 pp. 1–9.
- [33] Y. Song, G. Zhang, J. Lu, and H. Lu, "A fuzzy kernel c-means clustering model for handling concept drift in regression," in *2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2017. doi: 10.1109/FUZZ-IEEE.2017.8015515 pp. 1–6.
- [34] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, p. 1735–1780, nov 1997. doi: 10.1162/neco.1997.9.8.1735
- [35] L. Deng, "The mnist database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012. doi: 10.1109/MSP.2012.2211477
- [36] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009. [Online]. Available: <https://www.cs.utoronto.ca/~kriz/learning-features-2009-TR.pdf>
- [37] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 1126–1135. [Online]. Available: <https://proceedings.mlr.press/v70/finn17a.html>
- [38] Y. Chen, M. W. Hoffman, S. G. Colmenarejo, M. Denil, T. P. Lillicrap, M. Botvinick, and N. de Freitas, "Learning to learn without gradient descent by gradient descent," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017. doi: 10.5555/3305381.3305459 p. 748–756.
- [39] W. N. Street and Y. Kim, "A streaming ensemble algorithm (sea) for large-scale classification," in *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '01. New York, NY, USA: Association for Computing Machinery, 2001. doi: 10.1145/502512.502568. ISBN 158113391X p. 377–382.
- [40] J. Z. Kolter and M. A. Maloof, "Dynamic weighted majority: An ensemble method for drifting concepts," *The Journal of Machine Learning Research*, vol. 8, pp. 2755–2790, 2007. [Online]. Available: <http://jmlr.org/papers/v8/kolter07a.html>
- [41] M. P. S. Bhatia, "A two ensemble system to handle concept drifting data streams: recurring dynamic weighted majority," *International Journal of Machine Learning and Cybernetics*, vol. 10, 03 2019. doi: 10.1007/s13042-017-0738-9
- [42] A. Liu, J. Lu, and G. Zhang, "Diverse instance-weighting ensemble based on region drift disagreement for concept drift adaptation," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 293–307, 2020. doi: 10.1109/TNNLS.2020.2978523
- [43] B. Celik and J. Vanschoren, "Adaptation strategies for automated machine learning on evolving data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 9, p. 3067–3078, Sep. 2021. doi: 10.1109/tpami.2021.3062900
- [44] T. Chen, X. Chen, W. Chen, H. Heaton, J. Liu, Z. Wang, and W. Yin, "Learning to optimize: A primer and a benchmark," *Journal of Machine Learning Research*, vol. 23, no. 189, pp. 1–59, 2022. [Online]. Available: <http://jmlr.org/papers/v23/21-0308.html>
- [45] S. Wang, J. Sun, and Z. Xu, "Hyperadam: A learnable task-adaptive adam for network training," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 5297–5304, 07 2019. doi: 10.1609/aaai.v33i01.33015297
- [46] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016. doi: 10.48550/arXiv.1609.04747
- [47] T. Tieleman and G. Hinton, "Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude," *COURSERA: Neural Networks for Machine Learning*, vol. 4, pp. 26–31, 2012. [Online]. Available: <https://cir.nii.ac.jp/crid/1370017282431050757>
- [48] O. Wichrowska, N. Maheswaranathan, M. W. Hoffman, S. G. Colmenarejo, M. Denil, N. de Freitas, and J. Sohl-Dickstein, "Learned optimizers that scale and generalize," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017. doi: 10.5555/3305890.3306069 p. 3751–3760.
- [49] K. Lv, S. Jiang, and J. Li, "Learning gradient descent: better generalization and longer horizons," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017. doi: 10.5555/3305890.3305913 p. 2247–2255.
- [50] T. Chen, W. Zhang, Z. Jingyang, S. Chang, S. Liu, L. Amini, and Z. Wang, "Training stronger baselines for learning to optimize," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020. doi: 10.5555/3495724.3496339 pp. 7332–7343.
- [51] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '16. New York, NY, USA: ACM, 2016. doi: 10.1145/2939672.2939785. ISBN 978-1-4503-4232-2 pp. 785–794.
- [52] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, doi: 10.48550/arXiv.1603.04467, Software available from tensorflow.org.
- [53] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014. doi: 10.48550/arXiv.1412.6980

A Blockchain-based Transaction Verification Infrastructure in Public Transportation

Hidayet Burak Saritas^{1,2}

0000-0002-0425-3051

¹ Ege University International Computer Institute, Ege University International Computer Institute,

² Ege Teknopark, Ege University
Erzene Mah., Ankara Cad.,

35100 Bornova - Izmir, Türkiye

Email: burak.saritas@kentkart.com

Geylani Kardas

0000-0001-6975-305X

35100, Bornova – Izmir, Türkiye
Email: geylani.kardas@ege.edu.tr

Abstract—This paper proposes a new blockchain-based transaction verification infrastructure for co-payment and data verification for multi-modal public transportation systems. Our solution offers a decentralized platform that ensures secure co-payments and data integrity while addressing interoperability, data security and transactional transparency. With a private blockchain, transportation providers act as nodes and validated, consensus-approved transactions increase trust and transparency. A standardized data format and robust algorithms for data contribution by transport operators are developed as well as a model for operators, assets, and transactions. Including zero-knowledge proofs improves user privacy by allowing secure authentication without revealing sensitive data. We believe that this research may lead a closer collaboration between public transport operators and provide an enhanced user experience while enabling transport transaction security and data verification.

Index Terms—Blockchain Technology, Public Transportation, Co-payment Systems, Data Verification, Interoperability, Transaction Security, Decentralized Networks, Zero-Knowledge Proofs, Standardized Data Format, Privacy Preservation.

I. INTRODUCTION

WITH advancements in technology, the public transportation sector is expanding and offering a variety of options to a wider audience. In addition to the traditional methods, new alternatives like shared vehicles, scooters, and app-controlled taxis are becoming more popular [1]. Regulations in this field set limits to prevent unfair competition and encourage cooperation. Mobility partners have the freedom to make their own decisions. Various public transport providers need to work together on a common platform to manage this diversity and deliver an enhanced user experience. This can improve users' transportation experiences and reduce private vehicle usage. However, actors who offer different transportation methods have their own ticketing solutions. This may pose a significant challenge in the development process of a unified platform. These actors include ABT Kentkart [2], STIB-MIVB [3], MVV [4], and Whim [5]. Users need to adapt to various ticketing solutions, and each actor needs to

offer features such as payment, personalization, and usability [6] [7]. Therefore, creating a unified and accessible platform within the public transportation ecosystem is considered a significant innovation and challenge for the sector [8].

It is not easy to combine different actors of the public transportation industry. But it can be made simpler and also efficient with following solutions like a single application, account management and a single card. To combine services and share profits, researchers are constantly exploring easy integration methods to create a common language that all transportation solution providers can use. They are also working to establish a method for verifying every user transaction [8]. To solve these problems, this study proposes developing a blockchain-based solution that processes and verifies data produced by different transportation actors. Blockchain is a ledger built from computers in a distributed structure that cannot be controlled by any central authority. [9]. In essence, Distributed Ledger Technology (DLT) is a type of encrypted database that is distributed, shared, and serves as an irreversible and incorruptible information store [10]. It enables trust in transactions between two parties and eliminates the need for a central intermediary to provide this service. This allows for the secure transfer of unique assets, such as money, title deeds, and identification information, without intermediaries. Two users can conduct a financial transaction without the need for an intermediary institution [11]. They can communicate directly without any trust issues [12]. The paper aims to solve integration problems in the public transportation sector by enabling different operators and businesses to work together more effectively and reliably. It examines how blockchain technology's decentralized and reliable structure can increase transaction transparency and security in the sector.

The standard message package format was developed on a private blockchain network. This format creates an environment where transportation operators can add data. The study considered parameters such as speed and assets that validating operators must have. This way, different operators in the public transportation sector can use blockchain technology to in-

.Special acknowledgment is due to Kentkart A.Ş. Company for their generous support and guidance during the preparation of this paper.

tegrate with each other, trust each other, and query all transactions created with standard message format. Creating a trustworthy environment is crucial for this work. This can be achieved through a consensus mechanism and blockchain structure to confirm transactions. All businesses can join the blockchain network as validators, and data produced by any business can be added to the network after being approved by all nodes. The data added to the network in standard message format is trusted by everyone. To achieve this, a common data format and extensible approval mechanism have been created for the public transportation sector.

The rest of the paper is organized as follows: Section 2 reviews related work on blockchain applications in similar transportation sectors. In Section 3, we introduce the foundational elements and components of a blockchain environment optimized for public transportation. This section covers the architecture, roles, and interactions of transactions, as well as detailed descriptions of nodes, assets, and standardized message formats. In Section 4, we look at how to use decentralized technologies to keep people's information safe in public transportation. These technologies let people manage their digital identities without revealing sensitive information, which helps enhance people's privacy. We also look at how these technologies can be used with a single blockchain-based transaction verification system to verify user credentials. In Section 5, we summarize the contributions and implications of our blockchain-based infrastructure tailored for the public transportation sector.

II. RELATED WORK

Some notable research has been done on using blockchain for public transportation. Jayalath et al. [13] propose a micro-transaction model based on blockchain to improve service in Sri Lanka's public transportation sector. They focus on a ticketing system using an Ethereum-based blockchain to reduce transaction fees and improve service quality. This approach creates QR-based tickets for users to make micro-payments without third-party intermediaries. It reduces transaction costs.

Wang et al. [14] introduce "InterTrust," an interoperable blockchain architecture to enhance interoperability and reliability across various blockchain systems. The InterTrust model is for communication and interoperability among existing blockchain systems, which is a broader scope. However, our study aims at providing a specialized blockchain network for the public transportation domain.

Yang et al. [15] suggest a blockchain and Edge Computing-based communication system for maritime transportation. Their work uses blockchain and Edge Computing to improve Internet of Things (IoT) device performance and security in maritime environments. This is different from our work, which focuses on public transportation.

Enescu et al. [16] discuss a blockchain application to promote ecological transportation and reduce traffic congestion. They imagine a system where blockchain records transactions

and gives users digital currencies, which helps the environment and makes public transportation more popular.

Jabbar et al. [17] review blockchain applications in Intelligent Transportation Systems (ITS). They show how blockchain can improve transactional trust and efficiency. This supports various functionalities including automatic parking and fee payments.

Lastly, Chen et al. [18] describe a "Full-Spectrum Blockchain as a Service" (FSBaaS) approach with "Blockchain Lite" and "Hyperledger Fabric," focusing on providing blockchain services that cater to both centralized and decentralized frameworks. This study shows the flexibility needed in blockchain adoption.

Our study is different from the above mentioned noteworthy efforts because it suggests a way for public transport operators to trust each other. Each operator can check and verify data transactions without needing to ask anyone else. This makes sure that data is correct and that messages are the same across the network. It also makes easier to work together without having to rely on one person. Additionally, this study improves privacy and security in public transportation systems by using Decentralized Identifiers (DIDs), Verifiable Credentials (VCs), and Zero Knowledge Proofs (ZKPs) explained in later sections. It supports user privacy by authenticating users and transactions without exposing sensitive personal information. This approach not only secures digital identities and transactions but also sets a standard for using privacy-preserving technologies in public infrastructure. The proposed infrastructure helps advance blockchain technology and meets the needs of public transportation.

III. DEFINING BLOCKCHAIN MODEL ELEMENTS AND BASIC COMPONENTS

In this section, we introduce the main model elements and basic components of a blockchain-based transaction verification environment for public transportation.

A. High-Level Environment of Model Elements

In this study, we chose the Hyperledger Fabric blockchain environment and used the Raft consensus algorithm to establish a private blockchain environment. Hyperledger Fabric provides enhanced both control over transactions and privacy options. It allows for the creation of private channels, which limit visibility of specific transactions and data to certain network participants [19] [20]. Raft is known for efficiently handling high transaction volumes with low latency times [21]. This makes it ideal for verifying public transportation operations.

Figure 1 depicts the model elements operating in the high-level environment. These elements are defined as follows:

Transaction: Represents actions such as buying tickets or making payments, performed across various transportation modes including buses, cars, and scooters.

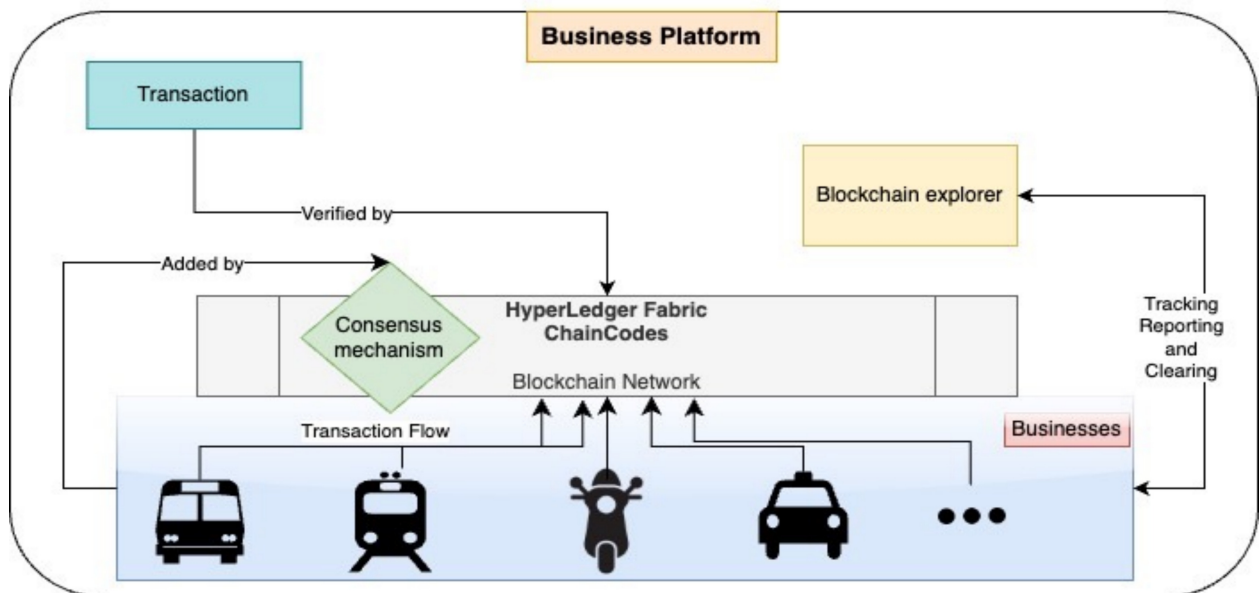


Fig. 1. High-Level System Design

Consensus Mechanism: Nodes (businesses) on the network verify transaction accuracy through a consensus mechanism [22]. This ensures that transactions are valid and their integrity is confirmed before being added to the blockchain.

Hyperledger Fabric Smart Contracts (Chaincodes): Transactions are processed using specific rules or smart contracts [19]. These smart contracts are utilized to automate, verify, and implement transaction logic.

Blockchain Network: Transactions take place and are recorded on the Hyperledger Fabric blockchain network.

Blockchain Explorer (Record Control Tool): An instrument utilized for visualizing and querying transactions, blocks, and other relevant data on the network. Users can track and report blockchain transactions with the help of this tool.

Businesses: Various enterprises serve as connected nodes on the network and participate in blockchain transactions.

The diagram in Figure 2 shows the defined classes for the fundamental components of the blockchain-based transaction verification system. The diagram demonstrates how a blockchain network is structured and how different components interact with each other. For example, a transaction executed by a node can trigger a smart contract, resulting in the addition of a block to the blockchain. The consensus mechanism verifies all processes and communicates using different data formats and messaging protocols. This plays a crucial role in maintaining the security and integrity of the blockchain network.

The “Node” class represents businesses and interacts with “SmartContract” and “Transaction”. The “Blockchain” and “Block” classes display the structural features of the

blockchain, while the “ConsensusMechanism” details how the nodes on the network reach a consensus. Additionally, the “DataFormat” and “MessagingProtocol” classes represent the data formats and communication protocols within the blockchain-based infrastructure.

B. Definitions and Contents of Model Elements

This section describes the model elements and their roles that form the basis of a private blockchain environment customized for public transportation. The structures used clearly demonstrate how nodes, entities, and transactions on the blockchain are identified, processed, verified, and integrated into the public transportation system. Examples and scenarios are provided to concretize the use of the standard message format in practice.

Nodes (Businesses): Each node in the blockchain network represents a business. Businesses manage various transportation services, such as buses, trams, and scooters, and have the authority to control and approve transactions on the network. Businesses act as 'nodes' to carry out their transactions and verify those of other businesses. They also contribute to the consensus mechanism to maintain the network's integrity and security.

Incorporating nodes into the network involves authentication and authorization processes using the Trusted Platform Module (TPM) [23]. The TPM is a hardware-based security module that provides node authentication. Each operating node contains a TPM chip. TPM securely generates and stores the node's private keys. The TPM module is activated by the business during installation and the first key pair is generated. TPM securely manages business identity and other cryptographic operations.

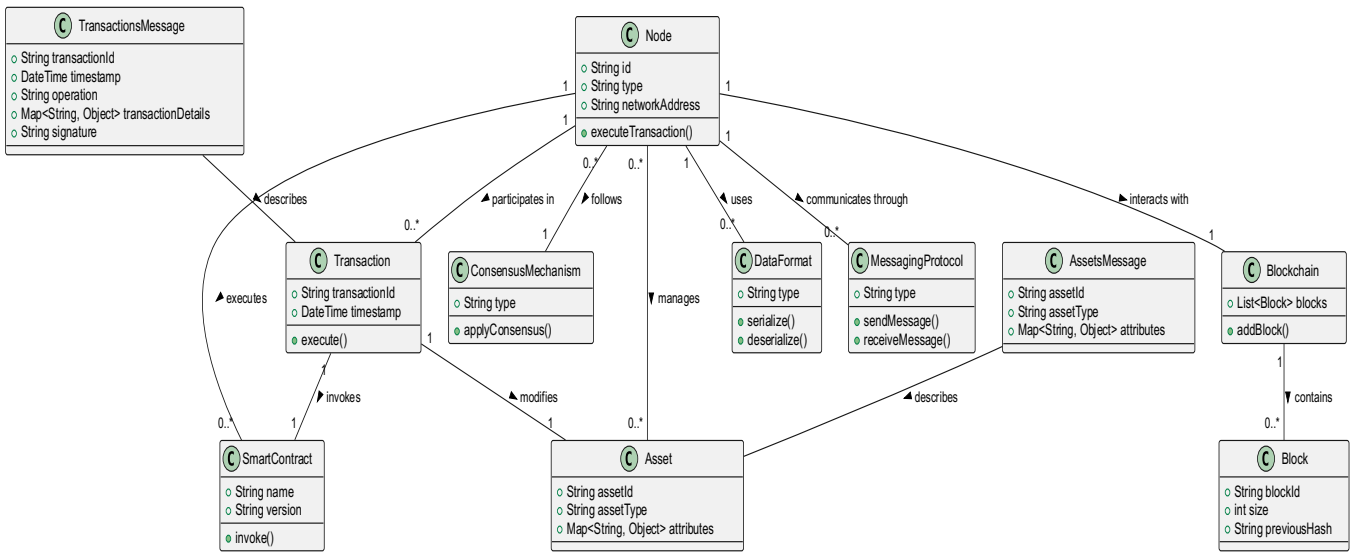


Fig. 2. Interactions between Components of Blockchain-based Transaction Verification System

Assets: Assets represent physical or digital items such as tickets, memberships, or payment records. For example, a ticket sales transaction creates an asset, while a boarding transaction represents the use of a ticket asset. An asset can be any item defined on the blockchain. Each entity has its own attributes, such as those given in the Table I defining the attributes of the Assets model element.

Each asset has an Asset ID and Asset Type, along with other predefined or optional attributes. Additional attributes may include time information, journey details, or any definition attributes necessary for verifying operations.

An extensible JSON structure for how assets are represented and used on the blockchain can be designed as shown in Listing 1. The initial design includes mandatory fields, and new fields can be added as needed.

Transactions: Transactions are defined as actions that create a change in the state of assets. Transaction types can include creating, updating, and deleting assets. Attributes of the Transactions model element are defined in Table II.

Transactions record all movements that users and businesses make on the network. Every transaction must be verified by other nodes in the network. Any unique transaction must contain the Transaction ID, which is a unique identifier, the timestamp of the transaction, and the type of transaction.

With the transaction examples given in Listing 2, it can be explained in detail how a transaction is initiated, how it progresses on the network and how it is concluded. For example, scenarios such as purchasing a ticket and boarding can be handled this way.

TABLE I.
MESSAGE STRUCTURE OF ASSETS MODEL ELEMENT

Attribute Name	Attribute Detail
Asset ID (assetId)	Each asset has a unique identifier
Asset Type (assetType)	The category into which the asset is classified (for example, 'ticket', 'membership')
Other Attributes	Other information that defines the properties of the asset (for example, validity period, price)

Listing 1. Message Structure of Assets Model Element

```

{
  "assetId": "123456",
  "assetType": "ticket",
  "attributes": {
    "issueDate": "2024-05-05 12:00:00",
    "expiryDate": "2024-06-05 12:00:00",
    "passengerId": "21412",
    "journeyDetails": {
      "origin": "Station A",
      "destination": "Station B",
      "departureTime": "2024-05-05 15:00:00"
    }
  }
}
    
```

TABLE II.
MESSAGE STRUCTURE OF TRANSACTIONS MODEL ELEMENT

Attribute Name	Attribute Detail
Transaction ID (transactionId)	Each transaction has a unique identifier
Timestamp	Indicates the time when the transaction took place
Operation Type	Indicates what type of action the operation is (for example, 'create', 'update', 'delete')

Listing 2. Message Structure of Transactions Model Element

```
{
  "transactionId": "tx123456789",
  "timestamp": "2024-05-05 12:00:00",
  "operation": "create",
  "transactionDetails": {
    "assetId": "123456",
    "assetType": "ticket",
    "attributes": {
      // Detailed information about the asset
    }
  },
  "signature": "DigitalSignatureOfTheUser"
}
```

Examples of transaction-specific data (payload) to be added to the message for different transactions in the standard message format to be sent over the network are given below. Listing 3 shows the ticket creation process, Listing 4 shows the membership update process, and Listing 5 shows sample "payload" information for fee payment. These payload samples represent the information required for various public transport operations and can be customized according to the needs of the transaction. Payload content may vary depending on the transaction type and the characteristics of the asset being processed.

C. Standardized General Message Format

It is necessary to determine a general message format to use all the message contents that we defined specifically for assets and transactions in the previous headings on the network. In this way, messages sent on the network can be standardized by using a common message format to produce transaction-specific data. Sample JSON structure and descriptions for the standard message format that can be used among all models in public transportation is defined as in Listing 6.

Listing 3. Payload of Ticket Creation Process

```
"payload": {
  "ticketNumber": "1234567890",
  "issueDate": "2024-01-01 10:00:00",
  "expiryDate": "2024-01-02 10:00:00",
  "passengerName": "Hakan Demir",
  "journeyDetails": {
    "origin": "Station A",
    "destination": "Station B",
    "departureTime": "2024-01-01
11:00:00"
  }
}
```

Listing 4. Payload of Membership Update Process

```
"payload": {
  "membershipId": "MEMB1234567",
  "memberName": "Hakan Demir",
  "validFrom": "2024-01-01",
  "validTo": "2024-01-01",
  "membershipType": "Gold",
  "additionalBenefits": ["Extra
Luggage", "Priority Boarding"]
}
```

Listing 5. Payload of Fee Payment

```
"payload": {
  "fareId": "FARE12345",
  "amountPaid": "15.00",
  "currency": "USD",
  "paymentMethod": "Credit Card",
  "transactionDate": "2024-01-01
12:30:00",
  "serviceType": "Tram"
}
```

This format is designed to encapsulate all necessary details for executing and verifying transactions, such as payments or asset transfers. It standardizes data for all parties involved in the transaction and maintains the system's integrity and reliability. Moreover, it supports various transportation modes and is versatile in different scenarios. The message's key components are the transaction type, transaction ID, timestamp, and invoked by, as well as transportation details. This part of the transportation message contains information about the mode of transport and route, including membership details for discounts or special fares.

Additionally, digital signature and consensus details are also present in general message format. Digital signature is used to authenticate the user's identity. The transaction's endorsing nodes, namely information, consensus timestamp, and consensus algorithm used are provided in the consensus detail part. Digital signatures and consensus details are pivotal for security, allowing for authenticated and verified transactions on the blockchain.

Listing 6. Standardized General Message Format

```

{
  "transactionType":
  "PaymentorAssetTransfer",
    "transactionId": "UniqueTransactionID",
    "timestamp": "2024-01-01 12:00:00",
    "invokedBy": "UserIDorBusinessID",
  "transportationDetails": {
    "modeOfTransport":
"bus/tram/scooter/minibus/metro",
    "routeId": "RouteID",
    "startLocation": "StartingLocation",
    "endLocation": "EndLocation",
  "fare": {
    "amount": "Amount",
    "currency": "Currency"
  },
  "membershipDetails": {
    "membershipId": "MembershipID",
    "validity": "MembershipValidity"
  }
},
"transactionDetails": {
  "assetId": "AssetID",
  "assetType": "AssetType",
  "operation": "create/update/delete",
  "payload": {
    // Customized data fields
    "journeyDetails": {
      "origin": "Station A",
      "destination": "Station B",
      "departureTime": "2024-05-05 15:00:00"
    },
    "fareDetails": {},
    "seatAllocation": {},
  }
},
  "signature": "DigitalSignatureOfTheUser",
  "consensusDetails": {
    "endorsedBy": ["NodeID1",
"NodeID2"],
    "consensusTimestamp": "2024-01-01
12:00:10",
    "consensusAlgorithm":
"ConsensusAlgorithmUsed"
  }
}

```

IV. ENHANCING USER PRIVACY IN PUBLIC TRANSPORTATION THROUGH DECENTRALIZED TECHNOLOGIES

This section explains how decentralized technologies can help protect user privacy in public transportation. It also looks at the challenges in digital systems that affect personal data security and how blockchain technology can help protect user

privacy. This study looks at how these technologies can be used to improve privacy in public transportation.

A. Privacy Challenges in the Digital Era

In today's world, it is crucial to securely store and process personal data due to the significant impact of social media and digitalization. Despite existing laws and regulations to protect user data, data breaches still occur, highlighting significant vulnerabilities. Storing personal data on company or institution servers is often the main cause of security breaches [24]. This is because central storage of personal data can be vulnerable to attacks and single point of failure can make it susceptible to cyber-attacks and unauthorized access. Personal data must be stored securely, and access should be restricted to authorized personnel only.

B. Emerging Solutions for Data Protection

To tackle the challenges addressed in Sect. IV.A, innovative technologies have been developed. These include the Decentralized Identifiers (DID) protocol [25], Verifiable Credentials (VC) [26], and Zero Knowledge Proof (ZKP) [27]. These technologies promote secure and digital storage of user data on individuals' devices, departing from traditional centralized systems. The DID protocol enables users to create and manage their digital identities without relying on central authorities. Verifiable Credentials enhance the dependability of digital identities and claims presented by users. They only include necessary data for a transaction. ZKP enables mathematical verification of information while maintaining privacy by not revealing personal details.

C. Application in Public Transportation Systems

A unified infrastructure allows users to securely access various public transportation services. The proposed study is different from conventional systems because users do not have to repeatedly share sensitive personal information with each service provider. Instead, it uses Verifiable Credentials stored in digital wallets, backed by the DID protocol. The claim information is stored in a data structure called the Sparse Merkle Tree [28]. This ensures data integrity by updating the root hash value, which is critical for verifying proofs. The root hash value is stored immutably on the blockchain.

D. Implementing a Privacy-Centric Approach

Users prove their eligibility to service providers by presenting ZKP-generated proofs alongside their credential information. These proofs can demonstrate eligibility as a student, teacher, elderly person, or person with a disability. Personal information is not required to be disclosed. Service providers verify these claims by referencing the proof and the root hash value on the blockchain. This model promotes a secure environment that minimizes personal data exposure and prioritizes privacy. Figure 3 shows how users can manage their digital identities securely and provide verification to service providers without revealing personal information. This is done by using technologies such as DID, VC, ZKP, and blockchain.

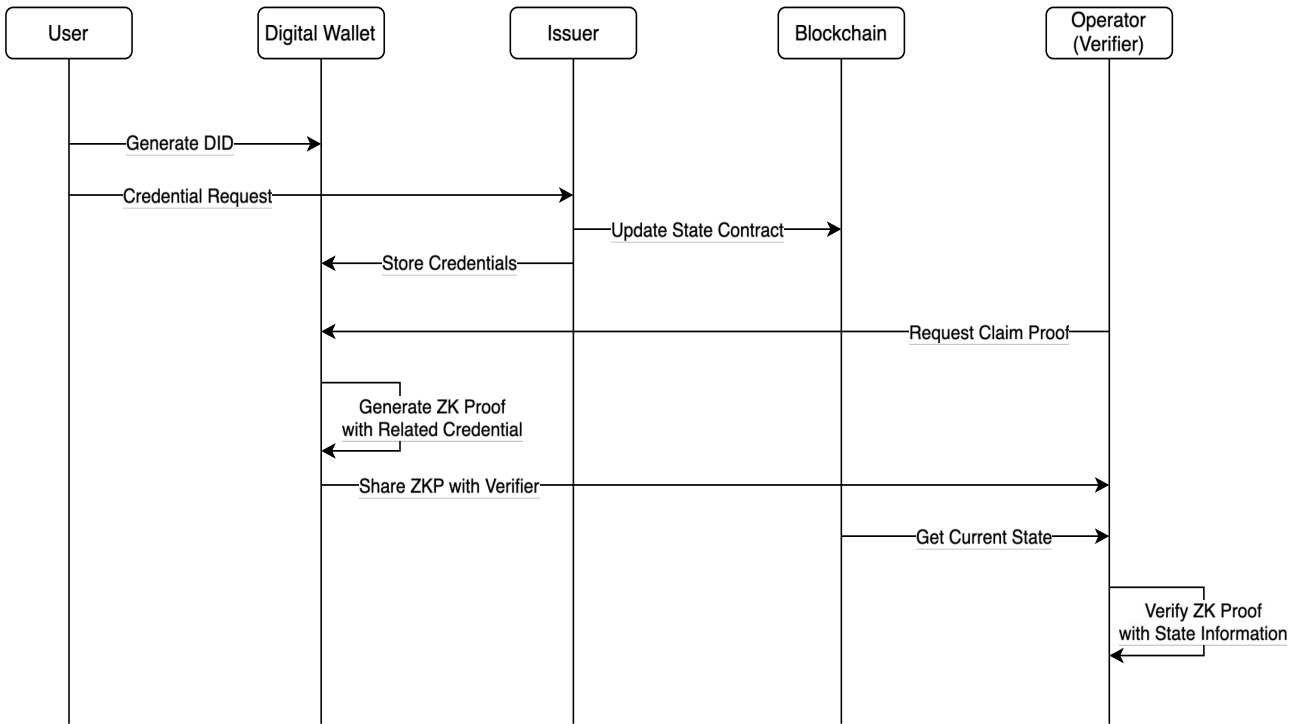


Fig. 3. User Privacy Protection in a Blockchain-Enabled Public Transportation System

Figure 3 shows how data and interactions flow between users, their digital wallets, issuers, the blockchain, and operators (verifiers) in a blockchain-based public transportation system. The system is designed to protect user privacy, and the user initiates the process by creating a Decentralized Identifier (DID) in their digital wallet. The issuer updates the relevant state contract on the blockchain and returns the approved credentials to the user's digital wallet for storage after requesting credentials.

The user's digital wallet generates a Zero-Knowledge Proof (ZKP) for the credential and shares it with the verifier, who is typically a transportation operator. The operator requests proof of claim and retrieves the current state from the blockchain to validate the transaction. The operator checks the ZKP using the state information from the blockchain. This process confirms the user's claim without compromising their privacy.

The verification flow ensures that the user's sensitive data remains secure. The Zero Knowledge Proof (ZKP) enables the user to demonstrate ownership of a valid credential without disclosing the credential itself. This use case shows how privacy-enhancing technologies can be used in public transportation systems. It utilizes blockchain technology to securely manage identities and verify transactions in a decentralized manner.

E. Outlook for Privacy Enhancement

The adoption of DID, VC, and ZKP represents a significant change in personal data protection and privacy enhancement. It is a positive development for the protection of personal data

and privacy. Service providers can conduct necessary verifications without compromising privacy, while users have greater control over their data. The study sets a new precedent for a digital ecosystem focused on security and user control [29].

V. CONCLUSION

A blockchain-based infrastructure for co-payment and data verification in the public transportation sector is presented in the paper. Our solution addresses interoperability, data security, and transactional transparency through a private blockchain. Transportation providers are nodes in this architecture and confirm transactions by consensus.

We analyze guarantee and reward systems to encourage participation and compliance across the network. Proposed infrastructure may also lead to implement systems carefully to establish trust among stakeholders and foster collaboration across multiple transportation services.

Design and deployment of smart contract APIs for network nodes are also discussed in the paper. These APIs are essential for running agreements autonomously, allowing fare settlements, service level agreements and more. They constitute the foundation of the system and enable smooth, scalable, and flexible services across the transportation network.

In addition, the paper addresses the strategic use of zero-knowledge proofs. This technology supports user privacy by authenticating users and transactions without revealing sensitive information. It meets today's digital privacy demands.

In conclusion, we believe that the proposed infrastructure is a significant step forward for public transport. It claims efficiency gains, enhanced security, and a user-centered approach. The system creates an environment for smart contract automation with advanced guarantee and reward mechanisms. This enables a seamless and secure public transportation experience. The presented research provides a roadmap for a harmonized, user-centric transportation network that respects privacy. This is especially relevant as urban mobility patterns change and require sophisticated solutions.

As the future work, we plan to develop and implement smart contracts for the public transportation system. These contracts will be tailored to the specific needs of the system. Additionally, the transaction verification mechanism will be operationalized within the blockchain environment. Transactions designed to guarantee user data confidentiality will be verifiable through the use of smart contracts.

REFERENCES

- [1] G. Oeschger, P. Carroll, and B. Caulfield, "Micromobility and public transport integration: The current state of knowledge," *Transportation Research Part D: Transport and Environment*, 2020. <https://doi.org/10.1016/j.trd.2020.102628>.
- [2] ABT Kentkart, "Automated fare collection system," Kentkart, 2022. <https://www.kentkart.com/solutions/automated-fare-collection-system>.
- [3] STIB-MIVB, "Ticket information," STIB-MIVB Ticket, 2022. https://www.stib-mivb.be/article.html?l=en&_guid=80bb5be7-429c-3810-a795-dfe836d62585.
- [4] MVV, "Online and handy ticket," MVV Ticketing, 2022. <https://www.mvv-muenchen.de/en/tickets-and-fares/online-und-handyticket/index.html>.
- [5] MaaS Global, "Whim - All your journeys," Whim, 2022. <https://whimapp.com>.
- [6] S. Kazi, M. Bagasrawala, F. Shaikh, and A. Sayyed, "Smart e-ticketing system for public transport bus," in *Proc. 2018 International Conference on Smart City and Emerging Technology (ICSCET)*, 2018, pp. 1-7. <https://doi.org/10.1109/ICSCET.2018.8537302>.
- [7] T. Khedekar, V. Jamdar, S. Waghmare, and M. L. Dhore, "FID automatic bus ticketing system," in *Proc. 2021 International Conference on Artificial Intelligence and Machine Vision (AIMV)*, 2021, pp. 1-6. <https://doi.org/10.1109/AIMV53313.2021.9670957>.
- [8] G. D. Pasquale, J. D. Bie, and J. Singh, "Ticketing in Mobility as a Service," International Association of Public Transport (UITP), 2022. <https://cms.uitp.org/wp/wp-content/uploads/2022/07/Report-Ticketing-MaaS-JULY2022-web.pdf>.
- [9] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," Technical Report, 2008. <https://bitcoin.org/bitcoin.pdf>.
- [10] H. Kakavand, N. Kost De Sevres, and B. Chilton, "The blockchain revolution: An analysis of regulation and technology related to distributed ledger technologies," *SSRN Electronic Journal*, 2016. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2849251.
- [11] I. Nath, "Data exchange platform to fight insurance fraud on blockchain," in *Proc. 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, 2016, pp. 821-825.
- [12] S. Gupta, S. Sinha, and B. Bhushan, "Emergence of blockchain technology: Fundamentals, working and its various implementations," *Proceedings of the International Conference on Innovative Computing & Communications (ICICC)*, 2020. <http://dx.doi.org/10.2139/ssrn.3569577>.
- [13] S. A. Jayalath, C. Rajapakse, and J. M. D. Senanayake, "A microtransaction model based on blockchain technology to improve service levels in the public transport sector in Sri Lanka," in *Proc. 2020 International Research Conference on Smart Computing and Systems Engineering (SCSE)*, 2020, pp. 82-89. <https://doi.org/10.1109/SCSE49731.2020.9313037>.
- [14] G. Wang and M. Nixon, "InterTrust: Towards an efficient blockchain interoperability architecture with trusted services," in *Proc. 2021 IEEE International Conference on Blockchain*, Melbourne, Australia, 2021, pp. 150-159. <https://doi.org/10.1109/Blockchain53845.2021.00029>.
- [15] T. Yang, Z. Cui, A. H. Alshehri, M. Wang, K. Gao, and K. Yu, "Distributed maritime transport communication system with reliability and safety based on blockchain and edge computing," in *IEEE Transactions on Intelligent Transportation Systems*. <https://doi.org/10.1109/ITITS.2022.3157858>.
- [16] F. M. Enescu, N. Bizon, G. Serban, and I. C. Hoarcă, "Environmental protection - Blockchain solutions for intelligent passenger transportation of persons," in *Proc. 2021 13th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)*, 2021, pp. 1-6. <https://doi.org/10.1109/ECAI52376.2021.9515026>.
- [17] Y. Jabbar, E. Dhib, A. B. Said, M. Krichen, N. Fetais, E. Zaidan, and K. Barkaoui, "Blockchain technology for intelligent transportation systems: A systematic literature review," in *IEEE Access*, vol. 10, 2022, pp. 20995-21031. <https://doi.org/10.1109/ACCESS.2022.3149958>.
- [18] Y. Chen, J. Gu, S. Chen, S. Huang, and X. S. Wang, "A full-spectrum blockchain-as-a-service for business collaboration," in *Proc. 2019 IEEE International Conference on Web Services (ICWS)*, 2019, pp. 219-223. <https://doi.org/10.1109/ICWS.2019.00045>.
- [19] Hyperledger Foundation, "Hyperledger Fabric Documentation," 2023. <https://hyperledger-fabric.readthedocs.io/>.
- [20] B. Reddy and P. S. Aithal, "Blockchain based service: A case study on IBM blockchain services & Hyperledger Fabric," *International Journal of Case Studies in Business, IT, and Education (IJCSBE)*, vol. 4, no. 1, May 2020, pp. 94-102. <https://ssrn.com/abstract=3611876>.
- [21] D. Ongaro, "In search of an understandable consensus algorithm (Extended Version)," Stanford University, 2014. <https://raft.github.io/raft.pdf>.
- [22] R. Awati, "Consensus algorithm," TechTarget. <https://www.techtarget.com/whatis/definition/consensus-algorithm>.
- [23] Trusted Computing Group, "TPM Main Specification," Trusted Computing Group, 2019. <https://trustedcomputinggroup.org/resource/tpm-library-specification/>.
- [24] Proceedings of the 17th Conference on Computer Science and Intelligence Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, D. Ślęzak (eds). ACSIS, Vol. 30, pages 685-694 (2022)
- [25] W3C, "Decentralized Identifiers (DIDs) v1.0," Jul. 19, 2022, W3C Recommendation. <https://www.w3.org/TR/did-core/>.
- [26] W3C, "Verifiable Credentials Data Model v2.0," World Wide Web Consortium. <https://www.w3.org/TR/vc-data-model/>.
- [27] Hyperledger Foundation, "Hyperledger AnonCreds: Anonymous Credentials with Zero-Knowledge Proofs," Hyperledger Wiki. <https://wiki.hyperledger.org/display/anoncreds>.
- [28] F. Haider, "Compact Sparse Merkle Trees," Oct. 6, 2018. <https://doi.org/10.31219/osf.io/8mcnh>.
- [29] A. Sherriff, K. Young, and M. Shea, "Editorial: Establishing Self Sovereign Identity with Blockchain," in *Front. Blockchain*, vol. 5, Art. no. 955868, Aug. 19, 2022. <https://doi.org/10.3389/fbloc.2022.955868>.

On Privacy of PRF+PUF-based Authentication

Ferucio Laurențiu Țiplea
0000-0001-6143-3641
Faculty of Computer Science
“Alexandru Ioan Cuza” University of Iași
Iași, Romania
Email: ferucio.tiplea@uaic.ro

Abstract—RFID-based authentication plays a crucial role in various fields, such as e-commerce, e-learning, e-business, healthcare, cloud, IoT, etc. At the same time, there is growing interest in using physically unclonable functions (PUFs) in RFID tags to protect against key corruption of pseudo-random functions (PRFs). In this paper, we discuss the privacy properties of PRF+PUF-based RFID authentication protocols in Vaudenay’s and the Hermans-Pashalidis-Vercauteren-Preneel (HPVP) models, considering two fundamental aspects: using temporary variables that might compromise privacy and using simulatable PUFs, a more realistic approach to ideal PUFs. Finally, we prove that a variant of a recently proposed RFID-based authentication protocol achieves strong privacy in the HPVP model.

I. INTRODUCTION

AUTHENTICATION is a process by which the validity of a particular assertion regarding an entity (entity authentication) or message (message or data origin authentication) is verified. Authentication plays a crucial role in information security because the authorization of access to information and data or the permission to carry out certain activities depends on it. Over time, many authentication techniques have been proposed, such as password-based authentication, certificate-based authentication, biometric authentication, token-based authentication, voice authentication, multi-factor authentication, and so on.

In this paper, we will look at authentication as a process by which one party, usually called the *verifier*, verifies the identity of another party, usually called the *prover*, by means of a protocol that takes place between the two parties. In some situations, the authentication process can include other parties, such as a server or a trusted authority. But, as we said, our paper focuses only on the authentication protocol developed between the prover and the verifier.

The basic requirements to be satisfied by an authentication process are security and privacy. In general, *security* means that no adversary can impersonate the prover or the verifier except with negligible probability. *Privacy* properties are much more nuanced and diverse, referring to anonymity, untraceability, unlinkability, etc.

Radio frequency identification (RFID) is a wireless communication technology between two parties, usually called *tag* and *reader*, through which the reader (playing the role of verifier) tries to uniquely identify the tag (which plays the role of prover). When the identification process is also completed by authentication, we speak of *RFID (RFID-based,*

RFID-enabled) authentication. RFID-based authentication is crucial in various fields such as e-commerce, healthcare, IoT, cloud, etc. [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18].

Over time, much effort has been dedicated to developing security and privacy models for RFID authentication. Two of the most important models are Vaudenay’s [19] and the Hermans-Pashalidis-Vercauteren-Preneel (HPVP) [20], [21] models. These models treat security identically. However, the privacy properties in Vaudenay’s model are stronger than those in the HPVP model. What is very important, however, is that these models propose a hierarchy of privacy properties and generic schemes for each level of privacy. The instantiation of such a scheme means nothing more than replacing the ideal cryptographic primitive, such as a pseudo-random function (PRF), pseudo-random generator (PRG), physically unclonable function (PUF), etc. with a practical construction. Suppose the said practical construction proves to be insecure at some time. In that case, it can be replaced with another practical construction, keeping the general scheme and the security and privacy results it enjoys.

Contributions and Paper Structure. One of the generic RFID schemes proposed in [19] is based on using a pseudo-random function. This scheme, generically called the *PRF-based RFID scheme*, ensures unilateral authentication and weak privacy both in Vaudenay’s and the HPVP models. Since using a PRF requires a key to be stored on the tag (on a prover’s device), the scheme cannot ensure privacy against adversaries who can corrupt the tag (since they will obtain the PRF key through corruption). To overcome this limitation, [22] endowed tags with PUFs for key storage. PUFs are typically assumed to be *physically unclonable* (it is infeasible to produce two PUFs that cannot be distinguished based on their challenge/response behavior), *unpredictable* (it is infeasible to predict the response to an unknown challenge), and *tamper-evident* (any attempt to physically access the PUF irreversibly changes the challenge/response behavior). In addition, PUFs are considered a less expensive alternative to non-volatile memory (NVM). The new construction based on PRF and PUF ensures unilateral authentication and destructive privacy in Vaudenay’s model, a higher level than weak privacy. However, there are two main issues with this scheme:

- Extending the scheme to ensure mutual authentication

raises the issue of how to use temporary variables;

- From a practical point of view, PUFs are nondeterministic. That means PUFs must be accompanied by an auxiliary mechanism to be used by the prover and verifier in the authentication protocol.

This paper addresses these two problems. We will discuss using temporary variables in the PRF+PUF-based RFID scheme in the Vaudenay and HPVP models. Then, we will simplify a recent authentication protocol based on PRFs and simulatable PUFs [23] and discuss the level of privacy achieved by it in these two models.

II. BASICS OF RFID SECURITY AND PRIVACY MODELS

This section will present some basic elements of Vaudenay's and the HPVP security and privacy models. We will be brief so the reader is referred to [19], [20], [21], [24], [25] for details. We use in our exposition *probabilistic polynomial time* (PPT) algorithms \mathcal{A} as defined in [25] that can consult *oracles*. For a set A , $a \leftarrow A$ means that a is uniformly at random chosen from A . If \mathcal{A} is a probabilistic algorithm, then $a \leftarrow \mathcal{A}$ means that a is an output of \mathcal{A} for some input.

The *authentication process* between a prover and a verifier requires the execution of a protocol between the two parties. When the communication between the prover and the verifier is carried out through radio waves, we say that we are dealing with an *RFID-based authentication*. In this context, the prover is seen as a resource-constrained small device, usually called a *tag*. However, the verifier called the *reader* is a computationally unrestricted powerful device that can perform any cryptographic operation. When the reader is near the tag, it transmits energy through radio waves, thus making the execution of an identification and authentication protocol possible. There are also scenarios where the tags have their own energy source.

Considerable effort has been put into the development of security and privacy models for RFID systems. Vaudenay's model [19] and the HPVP model [20], [21] are two of them, with major impact in the study of security and privacy properties of RFID systems. The HPVP model borrows the adversary model from Vaudenay's model, keeps the same approach to the security property, but treats privacy in a different way. If in Vaudenay's model privacy is based on indistinguishability between the RFID system instrumented by a challenger and the RFID system instrumented by a blinder (who does not know the secret elements in the system), the HPVP model treats privacy through indistinguishability between tags in the RFID system instrumented by a challenger. This second approach is closer to the security approach in the usual encryption systems. We will use these models throughout this work, so we recall their basic elements. First of all it is necessary to mention that the memory of a tag is typically split into *permanent* (or *internal*), used to store the state values of the tag, and *temporary* (or *volatile*), used to carry out the calculations required by the communication protocol. There are two types of temporary variables, *local*, used by tags only to do computations in a given protocol step, and *global*, that

get values in a given protocol step and are used in another protocol step.

Now, we can present Vaudenay's and the HPVP adversarial model. The oracles an adversary (PPT algorithm) can query in these models are those in the table in Figure 1.

The oracle *Corrupt* in Vaudenay's model returns only the current permanent state. We sometimes say that Vaudenay's model is *without temporary state disclosure* (TSD) [26], [27], [28], [29], [30], [31]. In *Vaudenay's model with TSD*, *Corrupt* returns the entire state of the tag, as in the HPVP model.

The adversaries can now be classified as follows:

- 1) (In both models) Adversaries with no access to *CreateInsider*. These are further classified according to the way the *Corrupt* oracle is used: *weak adversaries* (no access to *Corrupt*), *forward adversaries* (once they access the *Corrupt* oracle, the only oracle they can access is *Corrupt*), *destructive adversaries* (the tag is destroyed after corruption), *strong adversaries* (no restrictions), *narrow adversaries* (no access to *Result*);
- 2) (Only in the HPVP model) Adversaries with access to *CreateInsider*. The power of a destructive or strong adversary does not increase if he is given access to the *CreateInsider* oracle.

Security in Vaudenay's and the HPVP models means that no strong adversary has more than a negligible probability to make the reader authenticate an uncorrupted legitimate tag without having any tag authentication matching conversation. When the RFID scheme is with mutual authentication, besides the above requirement, it is asked that no strong adversary has more than a negligible probability to make an uncorrupted legitimate tag to authenticate the reader without having any reader authentication matching conversation.

Privacy generalizes well-known properties such as anonymity, unlinkability, untraceability, etc. It is treated differently in the two models. Vaudenay's model considers the *blinder* concept for a class \mathcal{C} of adversaries, which is a PPT algorithm \mathcal{B} that simulates the *Launch*, *SendReader*, *SendTag*, and *Result* oracles for adversaries in \mathcal{C} , without having access to the corresponding secrets. However, \mathcal{B} look passively at the communication between adversaries in \mathcal{C} and the other oracles allowed to it by the class \mathcal{C} (that is, \mathcal{B} gets exactly the same information as any adversary in \mathcal{C} when querying these oracles). The scheme is *C-private* in Vaudenay's model if no adversary in \mathcal{C} has more than a negligible advantage over $1/2$ to distinguish between protocol sessions in the real scheme from those in the scheme instrumented by some blinder.

An RFID scheme is considered *C-private* in the HPVP model if no adversary in \mathcal{C} can distinguish with more than a negligible probability over $1/2$ with which tag he communicated (the left or the right in the pairs drawn by him).

The previously defined adversary classes lead to a ranking of the privacy properties of RFID schemes as shown in the diagram in Figure 2.

Oracles in Vaudenay's model	Oracles in the HPVP model
<i>CreateReader</i> (\cdot): Unsupported	<i>CreateReader</i> (\cdot): Creates a new reader, and a unique reference R to it is returned
<i>CreateTag</i> ^{b} (ID): Creates a tag with the identifier ID . When $b = 1$, the tag is considered <i>legitimate</i> and registered in the server's database; otherwise ($b = 0$) it is considered <i>illegitimate</i> . A unique reference T to the tag is returned; the tag is considered <i>free</i>	<i>CreateTag</i> (ID): Creates a tag with the identifier ID and registers it in the server's database (that is, the oracle creates only legitimate tags). Duplicate tags with the same ID are accepted. A unique reference T to the tag is returned; the tag is considered <i>free</i>
<i>RegisterTag</i> (T, R): Unsupported	<i>RegisterTag</i> (T, R): Registers the tag T with the reader R
<i>Launch</i> (\cdot): Generates and outputs a new protocol session identifier π	<i>Launch</i> (R): Generates and outputs a new protocol session identifier π with the reader R
<i>DrawTag</i> (δ): The oracle chooses a number of free tags according to the distribution δ , let us say n , generates n temporary identities $vtag_1, \dots, vtag_n$, and outputs $(vtag_1, b_1, \dots, vtag_n, b_n)$, where b_i specifies whether the tag $vtag_i$ is legitimate or not. All these tags are considered now <i>drawn</i> . The oracle maintains a list Γ of drawn tags	<i>DrawTag</i> (T_0, T_1): Generates a fresh virtual tag reference $vtag$ that refers to either T_0 or T_1 , depending on the privacy game where the oracle is queried. The triple $(vtag, T_0, T_1)$ is included in a list Γ of drawn tags, and $vtag$ is returned by the oracle. The oracle returns \perp if one of the two tags is in the insider list, or one of the two tags is registered with a different set of readers than the other tag, or T_0 (T_1) is already referenced drawn
<i>Free</i> ($vtag$): Resets (erases) the temporary state of the tag referenced by $vtag$ and removes it from Γ	<i>Free</i> ($vtag$): Resets (erases) the temporary state of the tag referenced by $vtag$ and removes the corresponding triple from Γ
<i>SendTag</i> ($m, vtag$): Outputs the tag's answer when the message m is sent to the tag referred to by $vtag$. When m is the empty message, this oracle outputs the first message of the protocol instance π , assuming that the tag does the first step in the protocol	<i>SendTag</i> ($m, vtag$): Outputs the tag's answer when the message m is sent to the tag referred to by $vtag$. When m is the empty message, this oracle outputs the first message of the protocol instance π , assuming that the tag does the first step in the protocol
<i>SendReader</i> (m, π): Outputs the reader's answer when the message m is sent to it as part of the protocol instance π . When m is the empty message, abusively but suggestively denoted by \emptyset , this oracle outputs the first message of the protocol instance π , assuming that the reader does the first step in the protocol;	<i>SendReader</i> (R, m, π): Outputs the R 's reader answer when the message m is sent to it as part of the protocol instance π . When m is the empty message, abusively but suggestively denoted by \emptyset , this oracle outputs the first message of the protocol instance π , assuming that the reader does the first step in the protocol;
Outputs \perp if in session π the reader has not yet made a decision on tag authentication (this also includes the case when the session π does not exist), 1 if in session π the reader authenticated the tag, and 0 otherwise (this oracle is both for unilateral and mutual authentication)	<i>Result</i> (π): Outputs \perp if in session π the reader has not yet made a decision on tag authentication (this also includes the case when the session π does not exist), 1 if in session π the reader authenticated the tag, and 0 otherwise (this oracle is both for unilateral and mutual authentication)
<i>Corrupt</i> ($vtag$): Outputs the current permanent (internal) state of the tag referred to by $vtag$, when the tag is not involved in any computation of any protocol step (that is, the permanent state before or after a protocol step)	<i>Corrupt</i> (T): Outputs the current permanent and temporary state of the tag T . Remark that the corruption is with respect to a tag, not a virtual tag.
<i>CreateInsider</i> : Unsupported	<i>CreateInsider</i> (ID): Creates a tag and returns a unique reference T to it and its full state. The tag is included in a list of insider tags.

Fig. 1. Vaudenay's and the HPVP adversarial models




Fig. 2. Privacy levels in Vaudenay's and the HPVP model: "n-p" means "narrow p" and an arrow means "implication".

III. PUF-BASED RFID

Physically unclonable functions. One of the primary reasons that led to the development of *physically unclonable functions* (PUFs) was to find a method of protecting the secret keys against software and physical attacks [1]. A PUF can be considered a disordered physical system that can be challenged with external stimuli (challenges) to which it will react with corresponding responses. Unlike standard digital systems, the reaction of a PUF to a challenge depends on the micro- or nanoscale structural disorder of the PUF. Ideally, it is assumed that:

- 1) *Unclonability*: this disorder cannot be cloned or reproduced precisely, even by PUF's original manufacturer;
- 2) *Randomness*: the *response* r of the PUF to a *challenge* c is uniquely and uniformly at random chosen from the space of possible responses;
- 3) *Tamper-evident*: PUFs are tamper-evident (fully invasive attacks either damage or alter the functional behavior).

As a result, an *ideal PUF* defines a unique function P .

More realistically, a PUF is *noisy*. That is, it behaves like a non-deterministic function whose response depends on process variations, noise, environmental variables, and aging. However, it is assumed that all relevant environmental parameters are bounded, and the evaluation time of any given PUF has an upper bound. Therefore, two random evaluations of the PUF response given the same challenge might slightly vary with a Hamming distance between them bounded from above by a constant threshold. In such a case, one critical attribute of a *PUF* is the *reliability* of its responses, which estimates how consistently the responses can be generated against varying operating conditions.

Simulatable PUF. A *simulatable PUF* [1], [23] is a pair consisting of a noisy PUF and a parameterized model $SimPUF$ capable of computing a response r and its corresponding reliability confidence $conf$ in polynomial time for any given challenge c (i.e., $(r, conf) \leftarrow SimPUF(c)$), such that:

- 1) $SimPUF$ is constructed using one-time privileged access by an authorized party in a secure environment and subsequent acquisition of $SimPUF$ by any party is disabled;
- 2) if $r' \leftarrow PUF(c)$, then $P(r = r')$ is ϵ -close to 1 and $conf$ is ϵ -close to the reliability confidence of r' .

PUF tags. The (ideal) properties of PUFs mentioned above, as well as the technological progress aimed at achieving these properties, led to the proposal of security protocols that include them in various forms. Thus, we can mention protocols for oblivious transfer, bit commitment, key exchange, key generation, or authentication [1], [32], [13]. For example, [32] reviews key generators and authentication protocols based on PUFs proposed up to 2016. Among newer protocols, we mention [27], [23].

As far as we know, the first use of PUFs in RFID systems appears in [33], [34] to provide a solution to the problem of finding a private destructive RFID system in Vaudenay's model. Later, the use of PUFs in RFID systems gained momentum (see [28] for an ample discussion on this topic). The method of use is as follows. Tags are endowed with PUFs and store secret information (usually a secret key). When the tag authenticates itself to the reader, it interrogates the PUF and extracts that secret information, which will then be used in preparing the message for the reader.

Tags with PUFs embedded into them are usually called *PUF tags*. A *PUF-based RFID scheme* is an RFID scheme with PUF tags. The previously discussed adversary model is trivially extended to the case of PUF tags. We only need to discuss the *Corrupt* and *CreateInsider* oracles:

- 1) Due to PUF's tamper-evident property, no adversary with the possibility of corrupting PUF tags can obtain the secret information stored in the PUF. So, the $Corrupt(T)$ oracle returns only the full state of the tag;
- 2) Due to the non-clonability of PUFs, $CreateInsider(ID)$ creates a tag with the identity ID and lets the adversary simulate its PUF through a list of randomly generated pairs. This makes this oracle have a behavior similar to that of the original approach.

As a result of these, the classification and ranking of the privacy properties in Figure 2 remains the same for the case of PUF-based RFID schemes.

IV. PRF+PUF-BASED AUTHENTICATION

Obtaining an RFID mutual authentication scheme that would provide destructive privacy in Vaudenay's model was an open problem until 2010, when Sadeghi et al. [33], [34] managed to offer a solution. They started with the PRF-based RFID scheme and added PUF tags to generate PRF's

keys. Thus, corrupting the tags prevents the adversary from obtaining the PRF key. That is how the PRF+PUF paradigm was born, and it proved very useful in many constructions of authentication schemes later proposed.

In [35], a PRF+PUF-based RFID scheme was proposed, that achieves mutual authentication and destructive privacy in Vaudenay’s model (without TSD). The scheme is given in Figure 3. Here, ℓ_1 and ℓ_2 are of polynomial size in the security parameter λ , and $F = (F_K)_{K \in \mathcal{K}}$ is a PRF, where F_K is from $\{0, 1\}^{2\ell_1+1}$ to $\{0, 1\}^{\ell_2}$, for all $K \in \mathcal{K}_\lambda = \{K \in \mathcal{K} \mid |K| = \lambda\}$. Each tag is equipped with a (unique) PUF $P : \{0, 1\}^p \rightarrow \mathcal{K}_\lambda$ and has the capacity to compute F , where p is of polynomial size in λ . The internal state of the tag consists of a string $s \in \{0, 1\}^p$ randomly chosen as a seed to evaluate P . The reader maintains a database DB with entries for all legitimate tags. Each entry is a vector (ID, K) , where ID is the tag’s identity and $K = P(s)$, where P is the tag’s PUF.

The mutual authentication protocol is as follows. The reader sends initially a random x to the tag. On receiving it, the tag generates a random y , computes $K = P(s)$ and $z = F_K(0, x, y)$, erases K , and answers with (y, z) . The reader checks its database for a pair (ID, K) such that $z = F_K(0, x, y)$. If such a pair is found, it outputs ID ; otherwise, outputs \perp and randomly chooses $K \in \mathcal{K}_\lambda$. No matter of the two cases (K is found in the database or is randomly generated), the reader computes $w = F_K(1, x, y)$ and sends it to the tag. On receiving it, the tag computes $P(s)$ and $w' = F_{P(s)}(1, x, y)$. Finally, it outputs OK or \perp depending on the equality $w = w'$.

We notice that the tag erases the key K after using it in step 2. That prevents the key from being obtained through the tag’s TSD-corruption. As a result, the tag must recompute the key in step 4. However, this precaution does not guarantee that the protocol achieves destructive privacy in Vaudenay’s model with TSD. Let us consider the following narrow adversary \mathcal{A} against the scheme:

- 1) $CreateTag^1(ID)$;
- 2) $(vtag, 1) \leftarrow DrawTag(P(ID) = 1)$;
- 3) $\pi \leftarrow Launch()$;
- 4) $x \leftarrow SendReader(\emptyset, \pi)$;
- 5) $y, z \leftarrow SendTag(x, vtag)$;
- 6) $(s, x', y') \leftarrow CorruptTag(vtag)$;
- 7) If $y = y'$ then output 0 (the real privacy game) else output 1 (the blinded privacy game).

As one can see, \mathcal{A} creates a legitimate tag, draws it, runs a protocol session with the tag for the first two steps, and then corrupts the tag. If the temporary variable y' is not changed ($y' = y$), then the adversary plays the real privacy game with overwhelming probability. This is because the blinder does not know the tag’s internal state and, therefore, it cannot return the value of y , except with negligible probability. A similar attack can be mounted for the case of x .

As a conclusion, the scheme is not even narrow forward private in Vaudenay’s model with TSD.

In the HPVP model, the adversary cannot corrupt virtual tags; it can only corrupt physical tags. In addition, when a tag is released, its state is reset, which means that the adversary cannot obtain the values of the temporary variables after a protocol execution. As a result, we can simplify the protocol in Figure 3 without exposing it to corruption, as shown in Figure 4. Moreover, this new protocol ensures mutual authentication and strong privacy in the HPVP model (we will show this in a more general framework in the next section).

V. PRF+SIMPUF-BASED AUTHENTICATION

The use of a non-deterministic PUF P on a tag raises the problem of selecting the answer to a challenge c . However, having stored on the reader a SimPUF P' associated with P , there are procedures that can decide in polynomial time the answer of P on c . Such a procedure is TREVERSE proposed in [23]. This uses P' for possible responses of P to c , and the correct selection is made based on a pseudo-random function F . The authentication protocol is the one in Figure 5. The server initiates the protocol by sending a challenge c to the tag. The tag queries its (non-deterministic) PUF, obtains $r \leftarrow P(c)$, and responds with $(x_1, y = F_r(x_1))$, where x_1 is a random value and F is a pseudo-random function shared by the server and tag. When the server receives the tag response, it uses P' , the model of the PUF P , and the TRESERVE function to determine r^t , the possible response of P . The check is done by “ $y = F_{r^t}(x_1)$ ”. If such a value is found, the tag is authenticated and announced. Otherwise, the protocol is aborted. In the case of tag authentication, it sends the server a random value x_2 and receives $z = F_{r^t}(x_2)$. The value z is checked against $F_r(x_2)$. If the values match, the tag authenticates the server.

In [23], the authors presented a security analysis of the protocol in Figure 5. The adversarial model used is the following:

- The adversary can eavesdrop on the communication channel;
- The adversary can arbitrarily apply challenges via the publicly accessible interface to observe the tag’s response.

A supplementary assumption states that *SimPUF* enrollment is performed by the server in a secure environment using one-time privileged access, and such access is prohibited afterward.

Under these, the security analysis in [23] focused on brute force, replay, modeling, and physical attacks. However, from the protocol’s privacy point of view, [23] did not conduct any study. We note that the protocol uses r as a global temporary variable. Then, similarly to the previous section, an adversary can mount the following attack: corrupt the tag, get r , and verify the equality “ $z = F_r(x_2)$ ”. If it holds, the adversary plays the real privacy game; otherwise, the adversary plays the blinded privacy game. Therefore, the protocol cannot simultaneously ensure mutual authentication and narrow forward privacy in Vaudenay’s model.

Concerning privacy in the HPVP model, we will show below that the protocol achieves strong privacy. First, we

	Reader (DB)	Tag (P, s)
1	$x \leftarrow \{0, 1\}^{\ell_1}$	\xrightarrow{x}
2		$y \leftarrow \{0, 1\}^{\ell_1}, K = P(s)$ $z = F_K(0, x, y)$, erase K
3	If $\exists(ID, K) \in DB$ s.t. $z = F_K(0, x, y)$ then output ID (tag auth.) else output \perp , $K \leftarrow \mathcal{K}_\lambda$; $w = F_K(1, x, y)$	$\xleftarrow{y, z}$ \xrightarrow{w}
4		$K = P(s)$, $w' = F_K(1, x, y)$, erase K If $w = w'$ then output OK else output \perp

Fig. 3. PRF+PUF-based RFID scheme

	Reader (DB)	Tag (P, s)
1	$x \leftarrow \{0, 1\}^{\ell_1}$	\xrightarrow{x}
2		$y \leftarrow \{0, 1\}^{\ell_1}, K = P(s)$ $z = F_K(0, x, y)$
3	If $\exists(ID, K) \in DB$ s.t. $z = F_K(0, x, y)$ then output ID (tag auth.) else output \perp , $K \leftarrow \mathcal{K}_\lambda$; $w = F_K(1, x, y)$	$\xleftarrow{y, z}$ \xrightarrow{w}
4		$w' = F_K(1, x, y)$, erase K If $w = w'$ then output OK else output \perp

Fig. 4. A simplified variant of the PRF+PUF-based RFID scheme

simplify the protocol by eliminating unnecessary steps without changing its functionality. Figure 6 presents this new protocol.

Theorem 5.1: The mutual authentication scheme in Figure 6 provides strong privacy in the HPVP model, provided that P behaves randomly and F is a PRF.

Proof: Let Σ be the scheme in Figure 6. Assume that Σ is not strong private in the HPVP model, and let \mathcal{A} be a strong adversary that can break Σ 's privacy. We will show that there is an adversary \mathcal{B} that has a non-negligible advantage in the pseudo-randomness game with F . Let \mathcal{C} be a challenger for the pseudo-randomness game with F .

\mathcal{B} will simulate Σ (will be the challenger) in the privacy game that \mathcal{A} plays with Σ . So, \mathcal{B} will have to simulate the oracles for \mathcal{A} . \mathcal{B} does not know the secret parameters of the scheme but will want the simulation it performs to be indistinguishable from the real privacy game between \mathcal{A} and Σ . We will show below how the oracles are simulated:

- 1) \mathcal{B} keeps a list \mathcal{R} of readers that will be created by adversary, and a list of tags \mathcal{T}_R registered with each reader $R \in \mathcal{R}$. Initially, these lists are empty;
- 2) \mathcal{B} keeps a list \mathcal{T} of tags that will be created in the system in the order in which they are created. Each tag receives a fresh reference. We recall that the HPVP model allows the creation of several tags with the same identity. The

corrupted (insider) tags will be stored in a separate list $c\mathcal{T}$ ($i\mathcal{T}$), initially empty;

- 3) \mathcal{B} will simulate the tag T 's PUF as a list of challenge-response pairs. Initially, this list, denoted $P(T)$, is empty. When evaluating the PUF on c , \mathcal{B} looks in $P(T)$ a pair (c, r) , for some r . If such a pair is found, r will be returned as the value of P on c ; otherwise, \mathcal{B} generates a random value r , returns it, and includes (c, r) in $P(T)$;
- 4) \mathcal{B} keeps a list Γ of active triples $(vtag, T_0, T_1)$ as specified in the oracle $DrawnTag$. The oracle $Free(vtag)$ removes $(vtag, T_0, T_1)$ from Γ . Remark that Γ can contain at most one triple with $vtag$ in the first position;
- 5) \mathcal{B} will keep a list \mathcal{Q} of $(query, ext_answer)$ pairs, where $query$ is a query of \mathcal{A} and ext_answer is a possibly detailed information from which the answer to the query is extracted;
- 6) $CreateReader()$: \mathcal{B} generates a unique reader reference R , answers to \mathcal{A} with R , and includes R in \mathcal{R} . Moreover, $(CreateReader(), R)$ is included in \mathcal{Q} ;
- 7) $CreateTag(ID)$: \mathcal{B} generates a fresh tag reference T , associates it with ID , initializes $P(T)$ by the empty list, includes the pair (T, ID) in \mathcal{T} , and answers to \mathcal{A} with T . Moreover, $(CreateTag(ID), T)$ is included in \mathcal{Q} ;
- 8) $RegisterTag(T, R)$: \mathcal{B} includes T in the list \mathcal{T}_R and $(RegisterTag(T, R), \emptyset)$ in \mathcal{Q} ;

	Server (Reader) (SimPUF P' , PRF F)	Prover (Tag) (PUF P , PRF F)
1	$c \leftarrow \{0, 1\}^\ell$	\xrightarrow{c}
2		$x_1 \leftarrow \{0, 1\}^m, r \leftarrow P(c)$ $y := F_r(x_1)$ $\xleftarrow{x_1, y}$
3	If $fail \leftarrow TREVERSE(c, P', x_1, y)$ then <i>abort</i> else let r^t be its output (i.e., $y = F_{r^t}(x_1)$) authenticate tag	\xrightarrow{auth}
4		$\xleftarrow{x_2}$ $x_2 \leftarrow \{0, 1\}^m$
5	$z := F_{r^t}(x_2)$	\xrightarrow{z}
6		if $z = F_r(x_2)$ then <i>auth. server</i> else <i>abort</i>

Fig. 5. PRF + SimPUF-based authentication scheme in [23]

	Server (Reader) (SimPUF P' , PRF F)	Prover (Tag) (PUF P , PRF F)
1	$c \leftarrow \{0, 1\}^\ell$	\xrightarrow{c}
2		$r \leftarrow P(c), x_1 \leftarrow \{0, 1\}^m$ $y := F_r(x, 0)$ $\xleftarrow{x, y}$
3	If $fail \leftarrow TREVERSE(c, P', x, y)$ then <i>abort</i> else let r^t be its output (i.e., $y = F_{r^t}(x, 0)$) authenticate tag $z := F_{r^t}(x, 1)$	\xrightarrow{z}
4		if $z = F_r(x, 1)$ then <i>auth. server</i> else <i>abort</i>

Fig. 6. PRF + SimPUF-based strong private authentication scheme in the HPVP model

- 9) *Launch*(R): \mathcal{B} generates a fresh session identifier π , returns it to \mathcal{A} , and includes $(Launch(R), (R, \pi))$ in \mathcal{Q} ;
- 10) *DrawTag*(T_0, T_1): \mathcal{B} checks if the constraints of the *DrawTag* oracle are satisfied. If not, the answer is \perp . Otherwise, \mathcal{B} generates a fresh virtual tag reference $vtag$, includes $(vtag, T_0, T_1)$ in Γ , and answers with $vtag$. In \mathcal{Q} the pair $(DrawTag(T_0, T_1), \perp/vtag)$ is included, where $\perp/vtag$ is for the first/second case, resp.;
- 11) *Free*($vtag$): the triple whose first component is $vtag$ is removed from Γ (if it is in Γ). In this case, the pair $(Free(vtag), \emptyset)$ is included in \mathcal{Q} ;
- 12) *SendTag*($c, vtag$): \mathcal{B} extracts from Γ the triple whose first component is $vtag$. If no such triple exists, the answer is \perp . Otherwise, let $(vtag, T_0, T_1)$ be this triple. \mathcal{B} searches each list $P(T_0)$ and $P(T_1)$ for a pair with c in the first position. If one of the lists does not contain such a pair, \mathcal{B} generates a random r and includes (c, r) in that list. Now suppose that $(c, r_0) \in P(T_0)$ and $(c, r_1) \in P(T_1)$. \mathcal{B} randomly generates x , queries \mathcal{C} with $((r_0, x, 0), (r_1, x, 0))$ and returns (x, y) . In \mathcal{Q} ,

$(SendTag(c, vtag), \perp/(c, r_0, r_1, x, y))$ is included depending on one of the two cases above;

- 13) *SendReader*($R, (x, y), \pi$): Since x is generated randomly at each query of a tag, and y is calculated from x through a PRF function, x and y can be found in at most one tuple (c, r_0, r_1, x, y) previously computed by \mathcal{B} when R queried some tag by c . In addition, x and y can only appear independently with negligible probability. As a result, if \mathcal{B} does not find in \mathcal{Q} a tuple like the one above, it responds with \perp . Otherwise, it extracts the only tuple (c, r_0, r_1, x, y) , queries \mathcal{C} with $((r_0, x, 1), (r_1, x, 1))$, and returns the answer z of \mathcal{C} . In \mathcal{Q} , the pair

$(SendReader(R, (x, y), \pi), \perp/(c, r_0, r_1, x, y, z))$

is included depending on one of the two cases above;

- 14) *Result*(π): Having the entire history of the privacy game up to the moment of this query, \mathcal{B} can answer faithfully whether the tag was authenticated by the reader or not, or there is another case outside of these. In \mathcal{Q} , it will include $(Result(\pi), 1/0/\perp)$, depending on

the case;

- 15) *Corrupt*(T): The tag T has no permanent variables, its only global temporary variables being c and x . Through corruption its PUF is destroyed. As a result, if the tag is not in \mathcal{T} , \mathcal{B} has nothing to return to \mathcal{A} . Otherwise, it returns c and x , which \mathcal{A} has learned from previous communications anyway. \mathcal{B} moves then T from \mathcal{T} into the list $c\mathcal{T}$ of corrupted tags. The pair $(\text{Corrupt}(T), \emptyset/(c, x))$ is included in \mathcal{Q} (depending on the case);
- 16) *CreateInsider*(ID): \mathcal{B} creates a new tag reference T , associates it with ID , returns T to \mathcal{A} , and includes (T, ID) in $i\mathcal{T}$. Moreover, $(\text{CreateInsider}(ID), T)$ is included in \mathcal{Q} .

It is as clear as possible that the probability with which \mathcal{B} guesses to which component, left or right, \mathcal{C} applied the function F is precisely the probability with which \mathcal{A} guesses with which tag, left or right, played the privacy game for the Σ scheme. Therefore, the assumption that the protocol is not strongly private will contradict the pseudo-randomness of F . So, the protocol must be strongly private. ■

It is interesting to compare the protocol in Figure 6 with the one in Figure 4. The differences are only from the point of view of how the PUF is viewed, deterministic or non-deterministic.

VI. CONCLUSIONS

The privacy properties offered by RFID protocols are mainly studied using ad hoc techniques. That makes it that when we study privacy in reputable privacy models, such as the Vaudenay or HPVP model, we find that many RFID protocols do not ensure privacy at all [29], [30], [36], [37].

The present work makes a short foray into the PRF+PUF paradigm used in the last 15 years to construct RFID protocols. The emphasis falls on the use of temporary variables in the construction of such protocols, as well as on the difference in approach to PUFs as deterministic devices (such as the protocol in Figure 4) or non-deterministic (such as the protocol in Figure 6). The privacy study is conducted in each case using Vaudenay's and HPVP models.

The fact that the protocol in Figure 6 does not ensure narrow forward privacy in Vaudenay's model but ensures strong privacy in the HPVP model shows that the two models significantly differ in approach when we go beyond the forward level of privacy.

REFERENCES

- [1] U. Rührmair and M. van Dijk, "PUFs in security protocols: Attack models and security evaluations," in *2013 IEEE Symposium on Security and Privacy*, 2013, pp. 286–300.
- [2] W. Xie, L. Xie, C. Zhang, Q. Zhang, and C. Tang, "Cloud-based RFID authentication," in *2013 IEEE International Conference on RFID (RFID)*, 2013, pp. 168–175.
- [3] Z. Zhao, "A secure RFID authentication protocol for healthcare environments using elliptic curve cryptosystem," *J. Medical Syst.*, vol. 38, no. 5, p. 46, 2014.
- [4] C. Jin, C. Xu, X. Zhang, and J. Zhao, "A secure RFID mutual authentication protocol for healthcare environments using elliptic curve cryptography," *J. Medical Syst.*, vol. 39, no. 3, p. 24, 2015.
- [5] H. Xiao, A. A. Alshehri, and B. Christianson, "A cloud-based RFID authentication protocol with insecure communication channels," in *2016 IEEE Trustcom/BigDataSE/ISPA*, 2016, pp. 332–339.
- [6] H. Xu, J. Ding, P. Li, F. Zhu, and R. Wang, "A lightweight RFID mutual authentication protocol based on physical unclonable function," *Sensors*, vol. 18, no. 3, 2018.
- [7] M. Safkhani, Y. Bendavid, S. Rostampour, and N. Bagheri, "On designing lightweight RFID security protocols for medical IoT," *Cryptology ePrint Archive*, Paper 2019/851, 2019.
- [8] W. Liang, S. Xie, J. Long, K.-C. Li, D. Zhang, and K. Li, "A double PUF-based RFID identity authentication protocol in service-centric internet of things environments," *Information Sciences*, vol. 503, pp. 129–147, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020025519305857>
- [9] F. Zhu, P. Li, H. Xu, and R. Wang, "A lightweight RFID mutual authentication protocol with PUF," *Sensors*, vol. 19, no. 13, 2019. [Online]. Available: <https://www.mdpi.com/1424-8220/19/13/2957>
- [10] K. Fan, Q. Luo, K. Zhang, and Y. Yang, "Cloud-based lightweight secure RFID mutual authentication protocol in IoT," *Information Sciences*, vol. 527, pp. 329–340, 2020.
- [11] L. Xiao, H. Xu, F. Zhu, R. Wang, and P. Li, "SKINNY-based RFID lightweight authentication protocol," *Sensors*, vol. 20, no. 5, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/5/1366>
- [12] F. Zhu, P. Li, H. Xu, and R. Wang, "A novel lightweight authentication scheme for RFID-based healthcare systems," *Sensors*, vol. 20, no. 17, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/17/4846>
- [13] P. Gope and B. Sikdar, "A comparative study of design paradigms for PUF-based security protocols for iot devices: Current progress, challenges, and future expectation," *Computer*, vol. 54, no. 11, pp. 36–46, 2021.
- [14] M. Shariq, K. Singh, M. Y. Bajuri, A. A. Pantelous, A. Ahmadian, and M. Salimi, "A secure and reliable RFID authentication protocol using digital Schnorr cryptosystem for IoT-enabled healthcare in COVID-19 scenario," *Sustainable Cities and Society*, vol. 75, p. 103354, 2021.
- [15] V. Kumar, R. Kumar, S. Jangirala, S. Kumari, S. Kumar, and C.-M. Chen, "An enhanced RFID-based authentication protocol using PUF for vehicular cloud computing," *Security and Communication Networks*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:251239918>
- [16] M. Adeli, N. Bagheri, S. Sadeghi, and S. Kumari, "χperbp: a cloud-based lightweight mutual authentication protocol," *Peer Peer Netw. Appl.*, vol. 16, no. 4, pp. 1785–1802, 2023.
- [17] A. Kumar, K. Singh, M. Shariq, C. Lal, M. Conti, R. Amin, and S. A. Chaudhry, "An efficient and reliable ultralightweight RFID authentication scheme for healthcare systems," *Computer Communications*, vol. 205, pp. 147–157, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0140366423001329>
- [18] Y. Wang, R. Liu, T. Gao, F. Shu, X. Lei, G. Gui, and J. Wang, "A novel RFID authentication protocol based on a block-order-modulus variable matrix encryption algorithm," 2023.
- [19] S. Vaudenay, "On privacy models for RFID," in *Proceedings of the Advances in Cryptology 13th International Conference on Theory and Application of Cryptology and Information Security*, ser. ASIACRYPT'07. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 68–87.
- [20] J. Hermans, F. Pshalidis, Andreas and Vercauteren, and B. Preneel, "A new RFID privacy model," in *Computer Security – ESORICS 2011*, V. Atluri and C. Diaz, Eds. Berlin, Heidelberg: Springer Verlag, 2011, pp. 568–587.
- [21] J. Hermans, R. Peeters, and B. Preneel, "Proper RFID privacy: Model and protocols," *IEEE Transactions on Mobile Computing*, vol. 13, no. 12, pp. 2888–2902, Dec 2014.
- [22] F. Armknecht, A.-R. Sadeghi, I. Visconti, and C. Wachsmann, "On RFID privacy with mutual authentication and tag corruption," in *Proceedings of the 8th International Conference on Applied Cryptography and Network Security*, ser. ACNS'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 493–510.
- [23] Y. Gao, M. van Dijk, L. Xu, W. Yang, S. Nepal, and D. C. Ranasinghe, "TREVERSE: TRial-and-Error lightweight secure ReVERSE authentication with simulatable PUFs," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 1, pp. 419–437, 2022.
- [24] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, 3rd ed. Chapman & Hall/CRC, 2020.

- [25] M. Sipser, *Introduction to the Theory of Computation*. Cengage Learning, 2012.
- [26] F. L. Țiplea and C. Hristea, "Privacy and reader-first authentication in Vaudenay's RFID model with temporary state disclosure," *Cryptology ePrint Archive*, Report 2019/113, 2019, <https://eprint.iacr.org/2019/113>.
- [27] F. L. Țiplea and C. Hristea, "PUF protected variables: A solution to RFID security and privacy under corruption with temporary state disclosure," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 999–1013, 2021.
- [28] F. L. Țiplea, C. Andriesei, and C. Hristea, "Security and privacy of PUF-based RFID systems," in *Cryptography - Recent Advances and Future Developments*. IntechOpen, 2021, ISBN 978-1-83962-566-4.
- [29] F. L. Țiplea, "Lessons to be learned for a good design of private RFID schemes," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 4, pp. 2384–2395, 2022.
- [30] F. L. Țiplea, "Narrow privacy and desynchronization in Vaudenay's RFID model," *International Journal of Information Security*, vol. 22, pp. 563–575, June 2022.
- [31] F. L. Țiplea, C. Hristea, and R. Bulai, "Privacy and reader-first authentication in Vaudenay's RFID model with temporary state disclosure," *Comput. Sci. J. Moldova*, vol. 30, no. 3, pp. 335–359, 2022.
- [32] J. Delvaux, "Security analysis of PUF-based key generation and entity authentication," 2017.
- [33] A.-R. Sadeghi, I. Visconti, and C. Wachsmann, "PUF-enhanced RFID security and privacy," in *Workshop on secure component and system identification (SECSI)*, vol. 110, 2010.
- [34] —, *Enhancing RFID Security and Privacy by Physically Unclonable Functions*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 281–305.
- [35] C. Hristea and F. L. Țiplea, "Destructive privacy and mutual authentication in Vaudenay's RFID model," *Cryptology ePrint Archive*, Report 2019/073, 2019.
- [36] F. L. Țiplea, "On privacy of RFID-based authentication protocols," in *Proceedings of the 21st International Conference on Security and Cryptography - SECRIPT*, INSTICC. SciTePress, 2024, pp. 128–139.
- [37] F. L. Țiplea, "Security and privacy requirements for RFID schemes in healthcare: Case studies, solutions, and challenges," *Procedia Computer Science*, 2024, 28th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES 2024).

A Machine Learning Approach for Anxiety and Depression Prediction Using PROMIS[®] Questionnaires

Arthur R. S. Vitória*, Murilo O. Guimarães*, Daniel Fazzioni*, Aldo A. Díaz-Salazar*, Ana Laura S. A. Zara[†], Iwens G. Sene Junior*, Renato F. Bulcão-Neto*

*Institute of Informatics, Federal University of Goiás, Brazil

[†]Institute of Tropical Pathology and Public Health, Federal University of Goiás, Brazil

arthurvitoria@discente.ufg.br

0000-0003-1746-9668

Abstract—A mental disorder is a clinically significant disturbance in an individual’s cognition, emotional, or behavioral functioning. Mental disorders such as anxiety and depression can be accessed by psychiatrists using auxiliary tools such as the depression anxiety stress scale (DASS), patient reported outcome (PRO), patient reported outcome measures (PROMs) and patient reported outcomes measurement information system (PROMIS[®]). However, many individuals affected by the symptoms of mental disorders do not receive a proper diagnosis. In that context, this work proposes a machine learning approach to predict the score of anxiety and depression using PROMIS[®] questionnaires by performing a comparative study between supervised learning models to estimate the scores of anxiety and depression from individuals. Through the proposed model an average MAPE of 6.31%, R^2 of 0.76, and Spearman coefficient of 88.86 were achieved, outperforming widely used linear models such as support vector machines (SVM), random forest (RF), and gradient boosting (GB). In conclusion, the utilization of machine learning algorithms with PROMIS[®] questionnaires has shown promise as a methodology for assessing anxiety and depression scores from the participants’ perspective, aligning with their perceptions of well-being.

Index Terms—Anxiety, Depression, PROMIS[®], Machine Learning, Mental Health

I. INTRODUCTION

THE CONDITIONS manifested through the symptoms associated with mental disorders have a far-reaching impact that permeates and affects the personal relationships, occupational pursuits, and general well-being of millions of people [1]. The World Health Organization (WHO) has recorded a significant global incidence of mental disorders, with 970 million individuals affected, with anxiety and depression being the most common. Moreover, a notable surge in these numbers were observed during the COVID-19 pandemic, with an increase of 26% and 28% for anxiety and depression, respectively [2].

As reported by the Brazilian Ministry of Health (MH), 18.6 million Brazilians are affected by anxiety, and mental disorders constitute a significant contributing factor of disabilities across the Americas. In Latin America, Brazil is the country with the highest prevalence of depression, reaching around 15.5% of the population [3]. Additionally, symptoms of these mental

disorders between children and teenagers rose to 25.2% and 20.5% for depression and anxiety, respectively, during the COVID-19 pandemic. Underscoring the pressing need for comprehensive strategies aimed at enhancing the care and support for those affected by this reality.

Mental disorders can be assessed by psychiatrists using tools, such as DASS42 and DASS21 questionnaires, with a 42-item and 21-item, respectively. These self-administered instruments are designed to measure the magnitude of three negative emotional states: depression, anxiety and stress [4]. Mental disorders screening can also be carried out using patient reported outcome (PROs) and patient reported outcome measures (PROMs) with a validated accuracy. These measures rely on patient self-assessments covering aspects of well-being, including quality of life, symptom or symptom burden, experience of care, and mental health indicators like anxiety and depression [5], thus, providing insights from the patient’s perspective and ideas of their own health [6].

The treatment for these mental disorders usually consists of medications and psychotherapy. However, a substantial challenge arises from the delayed diagnosis of these mental disorders, resulting in limited access to timely and proper interventions [7]. As the condition deteriorates, the individual’s psychological capacity to seek treatment decreases, leading to an increase in the number of undiagnosed individuals [8]. Moreover, in primary care, only 50% of the patients with depression receive a diagnosis and only 15% receive a proper treatment [9]. In this way, emerging methodologies designed to enhance mental disorder screening are essential, as they can facilitate appropriate treatment and a decrease in the number of undiagnosed individuals.

Several studies have explored the application of statistical and machine learning models to predict therapy outcomes for a wide range of mental disorders [10]. Several studies primarily focus on making predictions regarding the long-term outcomes of patients with various conditions, either before diagnosis or during the course of treatment [11]. These conditions encompass schizophrenia [12], stress [13], [14], depression [15], [16], anxiety [17], [18], as well as other mental health disorders [19], [20]. However, since these mental disorders

often manifest gradually, with symptoms becoming discernible in their early stages, it is crucial to explore predictive modeling to anticipate and address potential issues before they advance further.

In that context, this work proposes a machine learning approach to predict the score of anxiety and depression using PROMs questionnaires by performing a comparative study between supervised learning models to estimate the scores of anxiety and depression from individuals. Furthermore, we investigated the relative contribution of the input variables to improve the understanding and importance of each variable related to the individuals' perspective and ideas of their own health.

This paper is organized as follows. In Section II the most relevant and related works are reported, describing the applied methodologies, results, and conclusions. In Section III, materials and methods used for the prediction of the scores of anxiety and depression are described, such as the proposed approach for training an MLP-based model and its evaluation. Section III-C describes and discusses the results of the relative importance of the input variables to the MLP model. Section V reports the results and Section VI describes the conclusions.

II. RELATED WORK

Literature search was conducted for articles that addressed the prediction of anxiety and depression levels using machine learning algorithms to provide context for current research and highlight advances and challenges in the prediction of mental health disorders. The databases used included PubMed, Scopus, Google Scholar, IEEE and IET. The keywords used were "anxiety", "depression", "machine learning", "deep learning", "questionnaires". As a result, 6 articles were selected as foundation for the current research.

In the research of [21], an attempt was made to determine five different levels of severity of anxiety, depression and stress. For the dataset, the DASS 21 questionnaire, which measures the level of anxiety, depression, and stress, was applied to 348 participants aged between 20 and 60. These questionnaires were used in five models: Decision Tree, Random Forests, Naïve Bayes, and k -Nearest Neighbor (KNN). Due to the unbalanced classes, the metric chosen was the F1-score. The model that showed the best F1-score for stress prediction was Random Forest with 71%. The best F1-score for depression was Naïve Bayes with 83%. For anxiety, none of the models performed well, reaching an average of 50%.

In the survey of [22], the DASS 42 questionnaire was used and filled in online by randomly chosen users between 2017 and 2019. Eight models were used: Naïve Bayes, Bayesian networks, k -nearest neighbors, multi-layer perceptron (MLP), radial basis function network (RBFN), random forest and J48. The results showed that the RBFN obtained the best accuracy in classifying the conditions of anxiety, depression and stress, with an average of 96% for each of the variables. A second test using the DASS 21 questionnaire showed that the MLP model achieved the best accuracy with 96% for stress, 93% for depression, and 98.8% for anxiety.

The study of [23] used a questionnaire with 55 questions. The answers of 604 participants were recorded. Depression was assessed using the Burns Depression Checklist (BDC). The study used various feature selection techniques to identify the most relevant ones. Six machine learning algorithms were applied to predict depression, with AdaBoost and the SelectKBest techniques achieving the highest accuracy of 92.56%.

In the study of [24], a group of people with autoimmune diseases were assessed, 637 participants completed a structured clinical interview for SDM-IV-TR axis disorders (SCID) and various PROMs. The models used include Logistic Regression (LR), Neural Networks (NN) and Random Forests, and were trained to predict anxiety and major depressive disorder (MDD) scores. As a result, the area under the curve (AUC) and Brier scores ranged from 0.87 to 0.91 and 0.07 (*i.e.*, no variation) for the MDD models and from 0.79 to 0.83 and 0.09 to 0.11 for the anxiety disorder models. In the LR and NN models, few PROMs items were needed to achieve an optimal performance.

The study of [25] aimed to develop an appropriate predictive model to diagnose anxiety and depression among elderly patients based on sociodemographic and health-related factors. Ten classifiers were evaluated using a dataset of 510 geriatric patients and tested using a 10-fold cross-validation method. The highest prediction accuracy of 89% was obtained with the random forest (RF) classifier. This RF model was tested with another dataset of 110 separate elderly patients for external validity. Its predictive accuracy was 91%, and the false positive rate was 10%, compared to the standard tool.

In [26], data collected from 935 university students in Bangladesh was used. The data included student demographic information, such as academic year, grade point average, and the results of two depression assessment scales: the Beck Depression Inventory (BDI-II) and the Anxiety, Depression and Stress Scales - Bangla Version (DASS 21-BV). In addition, the students answered a set of 16 questions related to the reasons for their depression. This data was used to train and test three machine learning algorithms: k -Nearest Neighbor (KNN), Random Forest (R), and Support Vector Machine (SVM). The RF algorithm showed the best performance, achieving an accuracy rate of 75% in identifying depression.

III. MATERIAL AND METHODS

This section presents a comprehensive overview of the data collection process, the definition of the multi-layer perceptron model (MLP), and how we measure the relative contribution of its input variables.

A. Dataset

A cohort study was carried out with primary data collection, approved by the Institutional Review Board (n. 5513411). The inclusion criteria were students, teachers or administrative technicians from a Brazilian public university, aged ≥ 18 years, attending one of the 22 monitored environments for at least two hours.

Data collection for this study was carried out using the Research Electronic Data Capture (REDCap) [27] platform, which is a highly versatile tool designed to meet the needs of research in the medical, public health and social sciences fields. Using REDCap, we developed customized electronic forms to collect relevant information.

At the beginning, participants filled in the PROMIS®Global-10 (or PROMIS®10) V2 Questionnaire [28] to identify the physical and mental health scores. After this initial phase, participants were instructed to take part in the follow-up process, supervised by recruiters. As some studies report, individuals mental health can be influenced by the quality of the air and the indoor environment in which they spend time [29], [30]. In that context, during this phase, after at least 30 minutes of stay in the monitored environment, the participants were invited to answer three sets of questions once a day:

- Perception of indoor air quality, temperature, and humidity [31].
- PROMIS® Anxiety Short Form: 8 questions related to the participants' feelings in order to calculate an anxiety score.
- PROMIS® Depression Short Form: 8 questions designed to calculate a depression score.

In the follow-up period, participants had the opportunity to answer the questionnaires up to 15 times. This made it possible to obtain a comprehensive and detailed longitudinal view of the trends and variations over time of the input variables under study.

The dataset has the following characteristics:

- Total variables: 80
- Date of first follow-up: 05/15/23
- Date of last follow-up: 08/14/23
- Total participants: 249
- Total tracking records: 1924

B. Artificial Neural Networks

An Artificial Neural Network (ANN) consists of fully connected artificial neurons organized into distinct layers: input, output, and hidden layers. The network's input is derived from samples, which provide the information for the ANN to assimilate and *learn* during its training. The most popular structure among ANNs is the feedforward [32]. In this work, we use a MLP, wherein information flows exclusively from the network's input layer to the output layer through a series of interconnected neurons.

Throughout the forward process during training an MLP model characterized by a predefined set of hidden layers (h_1, h_2, \dots, h_N) , each layer h_i containing n_i neurons, the sequence of calculations for each layer can be described as follows:

$$h_i^j = f \left(\sum_{k=1}^{n_{i-1}} w_{k,j}^{i-1} h_{i-1}^k \right) \quad (1)$$

where $i = 2, \dots, N$, since the first layer receives the input data directly, and $j = 1, 2, \dots, n_i$. The weights $w_{k,j}^i$ represents

the connection weights between the neuron k in the hidden layer i and the neuron j in the next hidden layer, the number of neurons in the i th hidden layer is denoted by n_i . After the forward pass, a cost function computes the prediction-reference difference, leading to error calculations. This error drives the back-propagation algorithm, adjusting MLP weights and biases. The process continues until it meets the predefined stopping criterion, which, in this study, was determined only by the number of epochs.

C. Feature Importance

Efforts to improve the interpretability of neural network models have led to the development of methods designed to estimate the relative contributions of input variables [33]. An ANN model can be represented as three set of layers, input, hidden, and output, and is often represented by weight matrices. However, the representation by weight matrices alone does not inherently convey valuable knowledge. In that context, these approaches often rely on the utilization of weight matrices within the neural network model to estimate the contributions of input variables [34].

The connection weight algorithm [35] measures the contributions of input variables by multiplying the raw connection weights from input neurons to hidden neurons with the connection weights from hidden neurons to output neurons and then summing these products across all the input neurons. The relative importance of a variable can be measured as

$$R_i = \sum_{j=1}^h w_{ij} w_{jo} \quad (2)$$

where R_i is the relative importance of the input variable i , h is the total number of hidden nodes in a layer, w_{ij} is the weight of the connection between input node i and hidden node j , and w_{jo} is the weight of the connection between the hidden node j and the output node o .

IV. EXPERIMENTS

The follow-up data of the participants contained a longitudinal view of the trends and variations over time of the input variables under study and were used in order to predict the scores of anxiety and depression. An overall structure of the experiments is presented in Figure 1, with a total of 249 participants and 74 input variables associated to each one. As part of the data preparation process, the collected responses were preprocessed carefully to enable model training. Additionally, the raw data is normalized between -1 and 1, allowing a adjust to a common magnitude scale, providing a more effective weight adjustment during the training time [36]. All the models were optimized using a grid search approach to identify hyperparameters that best matched the specific data settings with the aim of mitigating the risk of overfitting.

The participant data was partitioned into training, validation, and test sets using a 5-fold cross-validation approach. A representation of such process is depicted in Figure 2, each fold splits N participants into 5 sets for training, validation and

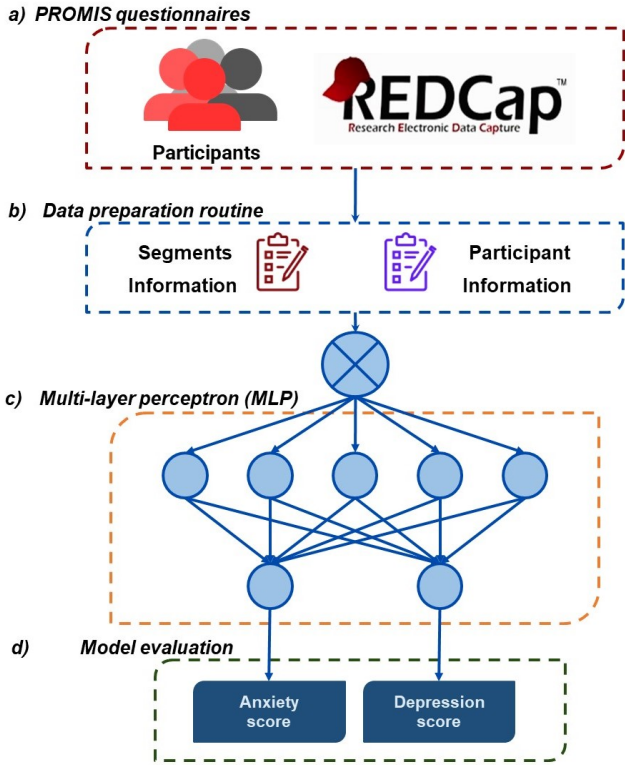


Fig. 1. Overall structure of the experiments. a) During the follow-up process, the input variables of the participants were collected using PROMIS[®] questionnaires along with questions related to their perception of indoor air quality; b) To prepare the data for the model training, we applied preprocessing routines that included various steps, such as data cleaning, feature engineering, and data normalization; c) An overall representation of the model structure, where each input variable is used as a feature, along with all participants information to train and test the model; d) A model evaluation routine was applied to systematically assess and analyze the outcomes achieved during training and testing phases.

testing. To prevent data leakage, all the follow-up information related to the participants was used exclusively within one set for each fold.

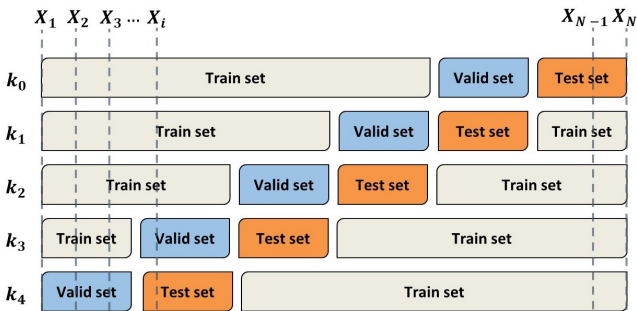


Fig. 2. The 5-fold cross-validation strategy.

The search space of hyperparameters in the proposed MLP model was tested extensively with the focus of determining the optimal number of neurons within each layer of the model, encompassing a wide array of combinations. The Adam [37]

algorithm was used as an optimizer for the training. The model's performance was assessed across multiple randomization seeds to measure their impact in different initialization. This assessment plays an important role in guiding the learning process away from undesired local minima, a critical aspect that needs careful consideration.

TABLE I
SEARCH SPACE OF MLP HYPERPARAMETERS.

Hyperparameter	Search Space
Hidden Layers	1, 2, 3
Number of Neurons	16, 32, 64, 128, 256, 512
Batch Size	128
Epochs	500
Seeds	73, 42, 10, 3407, 103
Learning rate	1e-4

Four other models were chosen as a baseline for comparison to the the MLP model: a model that combines multiple decision trees to make predictions (Random Forest); a supervised learning model that seeks to find a hyperplane for classification or regression (SVM); decision tree (DT), which makes decisions in a tree-like structure by splitting data based on rules; and gradient boosting (GB), an ensemble model that combines multiple DTs sequentially.

Each model was initialized and trained with the training set corresponding to the 5-fold cross-validation iteration. The training data included the input variables presented in the PROMIS[®] questionnaires, the perception of indoor air quality and, the corresponding output scores of anxiety and depression used as supervised labels. After training, each model was tested using the test set from the corresponding cross-validation iteration. This allowed the assessment of the model's performance on independent data.

Each model was assessed in terms of predictive capability using diverse criteria to enable the evaluation of various aspects of their generalization capabilities. These criteria include the mean absolute percentage error (MAPE), represented in Equation 3, where n represents the amount of data, \hat{y}_i is the predicted score and y_i is the measured score, addressing the performance of the models based on the average percentage difference between predicted and measured values.

$$MAPE = \frac{1}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{y_i} * 100 \quad (3)$$

The coefficient of determination R^2 described in Equation 4 expresses the capabilities of the model to accommodate the variance in the data. It shows the proportion of the variability in the dependent variable y that is explained by the independent variables, where n is the number of observations, y_i is the measured value, \hat{y}_i is the predicted value, and \bar{y} is the mean of measured values.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (4)$$

The Spearman Correlation Coefficient represented in Equation 5 allows a quantification of the degree and direction of the monotonic relationship between predicted and measured values, where n is the number of observation pairs and d_i represents the difference between ranked values of the two variables. It offers a non-parametric measure of the association of predicted and measured values. It varies from -1 (perfectly decreasing correlation) to 1 (perfectly increasing correlation), with 0 indicating no correlation.

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)} \quad (5)$$

V. RESULTS AND DISCUSSION

The dataset was split using the k -fold cross-validation, with $k = 5$ to divide the data in distinct folds, with each fold alternately serving as training, validation and test set. This allowed the models to be repeatedly trained and evaluated on different data combinations. In each iteration of k -fold, the models were trained on the training data, specific to that iteration, and their performance was assessed using the corresponding test data. The process was iterated on all possible fold combinations, resulting in a comprehensive evaluation of the models' performance across different scenarios. The results were evaluated in terms of MAPE, R^2 , and the Spearman Correlation Coefficient.




Fig. 3. Predictions for depression scores.

The difference in performance among the models is rooted in their intrinsic characteristics. Ensemble models, such as RF and GB, stand out for their capacity of diversification, variance reduction, and resist overfitting. They combine multiple decision trees, capturing different aspects of the data and aggregate predictions to provide robust results. In contrast, individual models like DT are simpler, with hierarchical rule structures that can limit their ability to capture complex relationships. They are prone to overfitting as its depths grows, especially in high-dimensional problems or data with intricate relationships.




Fig. 4. Predictions for anxiety scores.

Table II provides a summary of the predictive capabilities of anxiety and depression for different models. The results show a better performance of the RF and GB regressors among five selected models in predicting anxiety and depression. These two models exhibited relatively lower MAPE values and higher R^2 scores, demonstrating strong ability to explain the variance in the data. In contrast, the DT model demonstrated an acceptable performance in predicting anxiety, but encountered challenges in explaining the variation in depression data, as evidenced by its lower R^2 score. The RF model displayed a robust performance, yielding a relatively low MAPE of 5.51% for anxiety and 9.16% for depression. Moreover, it achieved a high R^2 score, indicating its strong capacity to elucidate the variance within the data. The SVM model delivered reasonable results with a MAPE of 8.10% for anxiety and 9.10% for depression, although a slightly lower R^2 value when compared to the RF. The DT grappled with predicting depression, which is evident in its higher MAPE of 10.37% and a lower R^2 of 0.30. In contrast, the GB model achieved a MAPE of 7.05% for anxiety and 9.81% for depression, while maintaining fairly low R^2 values.

Figure 3 and Figure 4 show the predictions of the best trained model using k -fold and a grid-search optimization for anxiety and depression scores. A visual inspection in these results shows that the proposed MLP model performed well in both scenarios, where the only significant discrepancies that can be point out are in segments that reach the minimum score due to the lack of information provided by the participant. Using the connection weights algorithm proposed by [35], the relative contribution of each input variable was accessed and showed in Figure 5. It is evident that throughout the training process, the model consistently exhibited a tendency to prioritize similar variables for both anxiety and depression. Moreover, these input are identified in Table III and Table IV, where they are sorted in a descending order, where the first feature is the most relevant. It is worth mention that despite seventy four input variables, the model learned to give

TABLE II
RESULTS OF MAPE, R^2 , AND CORRELATION COEFFICIENT OF SPEARMAN FOR ANXIETY AND DEPRESSION.

Models**	Anxiety			Depression		
	MAPE	R^2	Spearman	MAPE	R^2	Spearman
RF	5.51% ± 0.02	0.79 ± 0.16	0.86 ± 0.10	9.16% ± 0.02	0.54 ± 0.02	0.82 ± 0.08
SVM	8.10% ± 0.01	0.66 ± 0.08	0.84 ± 0.05	9.10% ± 0.01	0.60 ± 0.06	0.82 ± 0.05
DT	5.67% ± 0.03	0.62 ± 0.31	0.80 ± 0.15	10.37% ± 0.03	0.30 ± 0.06	0.75 ± 0.13
GB	7.05% ± 0.01	0.76 ± 0.05	0.86 ± 0.03	9.81% ± 0.01	0.53 ± 0.06	0.83 ± 0.13
MLP *	6.98% ± 0.02	0.72 ± 0.03	0.85 ± 0.01	5.64% ± 0.04	0.80 ± 0.02	0.91 ± 0.02

** Best results for optimization and cross validation processes.

TABLE III
INPUT RELATIVE CONTRIBUTION FOR ANXIETY SCORES.

Feature Importance – Anxiety Prediction

How often have you been bothered by emotional problems, such as feeling anxious, depressed or irritable?

General clinical symptoms (weakness, tiredness, nausea, others).

Neurological symptoms (headache, migraine, dizziness, others).

Respiratory symptoms (sneezing, stuffy nose, runny nose, difficulty breathing, dry throat or sore throat, others).

In general, how would you rate your mental health, including your mood and your ability to think?

Dermatological symptoms (burning skin, redness, allergies, etc.)

In general, rate how well you manage to carry out your frequent social activities and functions

(including activities at home, at work and in the community,
and responsibilities as a parent, child, spouse, employee, friend, etc.).

Do you carry out any professional activities?

Do you have any of these chronic diseases? (choice = sinusitis)

How well can you perform daily physical activities such as walking, climbing stairs, carrying groceries from the supermarket or moving a chair?

TABLE IV
INPUT RELATIVE CONTRIBUTION FOR DEPRESSION SCORES.

Feature Importance – Depression Prediction

How often have you been bothered by emotional problems, such as feeling anxious, depressed, or angry?

General clinical symptoms (weakness, tiredness, nausea, others).

Neurological symptoms (headache, migraine, dizziness, others).

In general, how would you rate your mental health, including your mood and your ability to think?

Respiratory symptoms (sneezing, stuffy nose, runny nose, difficulty breathing, dry throat or sore throat, others).

Dermatological symptoms (burning skin, redness, allergies, etc.).

In general, rate how well you are able to carry out your frequent social activities and functions (including activities at home, at work and in the community, and responsibilities as

a parent, child, spouse, employee, friend, etc.).

Do you have any of these chronic diseases? (choice = sinusitis)

On average, how would you rate your tiredness?

Do you carry out any professional activities?




Fig. 5. Relative importance of the input variables for predict the anxiety and depression scores.

more importance around questions related to the general well-being of participants, as well as questions about symptoms they presented during the follow-up process, including general clinical, respiratory, and neurological, and dermatological in case of anxiety. Furthermore, some questions concerning participants' views on physical health, including aspects like fatigue, engagement in social activities, and functions that encompasses personal activities can be observed as more important.

VI. CONCLUSION

This work investigates the potential of self-assessments covering aspects of well-being, mental health indicators, and the perception of indoor air quality, collected through PROMIS® questionnaires to measure scores of anxiety and depression for several participants. The dataset encompasses data from 219 participants who volunteered during a follow-up period. Comprehensive information regarding well-being status and perception of the surroundings were periodically collected through questionnaires in predefined environments, with a total of 1924 tracking records. However, despite a good amount of data from the participants, there was a lack of more representative scores for severe depression and anxiety to achieve a model with a good generalization.

A comparative study of supervised learning methods was conducted to measure scores of anxiety and depression for all the questionnaires under evaluation. Although simpler models demonstrated exceptional performance for anxiety scores, achieving 5.51% of MAPE and 0.79 of R^2 for the RF model, the prediction performance for depression scores was not as impressive, with the lowest MAPE of 9.16% and a highest R^2

of 0.54. The proposed MLP model outperformed the baseline, achieving an average MAPE for both anxiety and depression of 6.31%, an R^2 of 0.76, and a Spearman coefficient of 88.86. This suggests that machine learning models with the ability of capture more complex patterns in data, such as MLPs, might be better suited for addressing the prediction of anxiety and depression scores. Through the connection weight algorithm it was possible to determine the relative importance of each input feature. Additionally, future work could include other mental disorders, such as burnout [38].

REFERENCES

- [1] A. Priya, S. Garg, and N. P. Tigga, "Predicting anxiety, depression and stress in modern life using machine learning algorithms," *Procedia Computer Science*, vol. 167, pp. 1258–1267, 2020.
- [2] "Mental health and covid-19: Early evidence of the pandemic's impact," 2022.
- [3] MH, "Depressão," <https://www.gov.br/saude/pt-br/assuntos/saude-de-a-a-z/d/depressao>, 2023, last accessed 05 Oct 2023.
- [4] L. Parkitny and J. McAuley, "The depression anxiety stress scale (dass)," *Journal of physiotherapy*, vol. 56, no. 3, p. 204, 2010.
- [5] T. Weldring and S. M. Smith, "Article commentary: patient-reported outcomes (pros) and patient-reported outcome measures (proms)," *Health services insights*, vol. 6, pp. HSI-S11 093, 2013.
- [6] L. G. Tennenhouse, R. A. Marrie, C. N. Bernstein, L. M. Lix *et al.*, "Machine-learning models for depression and anxiety in individuals with immune-mediated inflammatory disease," *Journal of psychosomatic research*, vol. 134, p. 110126, 2020.
- [7] R. Razavi, A. Gharipour, and M. Gharipour, "Depression screening using mobile phone usage metadata: a machine learning approach," *Journal of the American Medical Informatics Association*, vol. 27, no. 4, pp. 522–530, 2020.
- [8] A. L. B. Melo and A. L. F. Alves, "Aplicação de técnicas de aprendizagem de máquina para diagnóstico de depressão, ansiedade e estresse," in *Anais do IX Encontro Nacional de Computação dos Institutos Federais*. SBC, 2022, pp. 13–16.

- [9] A. J. Mitchell, A. Vaze, and S. Rao, "Clinical diagnosis of depression in primary care: a meta-analysis," *The Lancet*, vol. 374, no. 9690, pp. 609–619, 2009.
- [10] S. Hornstein, V. Forman-Hoffman, A. Nazander, K. Ranta, and K. Hilbert, "Predicting therapy outcome in a digital mental health intervention for depression and anxiety: A machine learning approach," *Digital Health*, vol. 7, p. 20552076211060659, 2021.
- [11] A. B. Shatte, D. M. Hutchinson, and S. J. Teague, "Machine learning in mental health: a scoping review of methods and applications," *Psychological medicine*, vol. 49, no. 9, pp. 1426–1448, 2019.
- [12] N. C. Hettige, T. B. Nguyen, C. Yuan, T. Rajakulendran, J. Baddour, N. Bhagwat, A. Bani-Fatemi, A. N. Voineskos, M. M. Chakravarty, and V. De Luca, "Classification of suicide attempters in schizophrenia using sociocultural and clinical features: A machine learning approach," *General hospital psychiatry*, vol. 47, pp. 20–28, 2017.
- [13] J. L. Hagad, K. Moriyama, K. Fukui, and M. Numao, "Modeling work stress using heart rate and stress coping profiles," in *Principles and Practice of Multi-Agent Systems: International Workshops: IWEC 2014, Gold Coast, QLD, Australia, December 1-5, 2014, and CMNA XV and IWEC 2015, Bertinoro, Italy, October 26, 2015, Revised Selected Papers 5*. Springer, 2016, pp. 108–118.
- [14] Y. Nakashima, J. Kim, S. Flutura, A. Seiderer, and E. André, "Stress recognition in daily work," in *Pervasive Computing Paradigms for Mental Health: 5th International Conference, MindCare 2015, Milan, Italy, September 24-25, 2015, Revised Selected Papers 5*. Springer, 2016, pp. 23–33.
- [15] T. Hajek, C. Cooke, M. Kopecek, T. Novak, C. Hoschl, and M. Alda, "Using structural mri to identify individuals at genetic risk for bipolar disorders: a 2-cohort, machine learning study," *Journal of Psychiatry and Neuroscience*, vol. 40, no. 5, pp. 316–324, 2015.
- [16] Y. Hou, J. Xu, Y. Huang, and X. Ma, "A big data application to predict depression in the university based on the reading habits," in *2016 3rd International Conference on Systems and Informatics (ICSAI)*. IEEE, 2016, pp. 1085–1089.
- [17] F. Liu, W. Guo, J.-P. Fouché, Y. Wang, W. Wang, J. Ding, L. Zeng, C. Qiu, Q. Gong, W. Zhang *et al.*, "Multivariate classification of social anxiety disorder using whole brain functional connectivity," *Brain Structure and Function*, vol. 220, pp. 101–115, 2015.
- [18] T. Tran and R. Kavuluru, "Predicting mental conditions based on "history of present illness" in psychiatric notes with deep neural networks," *Journal of biomedical informatics*, vol. 75, pp. S138–S148, 2017.
- [19] M. Khondoker, R. Dobson, C. Skirrow, A. Simmons, and D. Stahl, "A comparison of machine learning methods for classification using simulation with multiple real data examples from mental health studies," *Statistical methods in medical research*, vol. 25, no. 5, pp. 1804–1823, 2016.
- [20] W. J. Bosl, T. Loddenkemper, and C. A. Nelson, "Nonlinear eeg biomarker profiles for autism and absence epilepsy," *Neuropsychiatric Electrophysiology*, vol. 3, no. 1, pp. 1–22, 2017.
- [21] A. Priya, S. Garg, and N. P. Tigga, "Predicting anxiety, depression and stress in modern life using machine learning algorithms," *Procedia Computer Science*, vol. 167, pp. 1258–1267, 2020.
- [22] P. Kumar, S. Garg, and A. Garg, "Assessment of anxiety, depression and stress using machine learning models," *Procedia Computer Science*, vol. 171, pp. 1989–1998, 2020, third International Conference on Computing and Network Communications (CoCoNet'19).
- [23] M. S. Zulfiker, N. Kabir, A. A. Biswas, T. Nazneen, and M. S. Uddin, "An in-depth analysis of machine learning approaches to predict depression," *Current Research in Behavioral Sciences*, vol. 2, p. 100044, 2021.
- [24] L. G. Tennenhouse, R. A. Marrie, C. N. Bernstein, and L. M. Lix, "Machine-learning models for depression and anxiety in individuals with immune-mediated inflammatory disease," *Journal of Psychosomatic Research*, vol. 134, p. 110126, 2020.
- [25] A. Sau and I. Bhakta, "Predicting anxiety and depression in elderly patients using machine learning technology," *Healthcare Technology Letters*, vol. 4, no. 6, pp. 238–243, 2017.
- [26] A. A. Choudhury, M. R. H. Khan, N. Z. Nahim, S. R. Tulon, S. Islam, and A. Chakrabarty, "Predicting depression in bangladeshi undergraduates using machine learning," in *2019 IEEE Region 10 Symposium (TENSymp)*, 2019, pp. 789–794.
- [27] P. A. Harris, R. Taylor, R. Thielke, J. Payne, N. Gonzalez, and J. G. Conde, "Research electronic data capture (redcap)—a metadata-driven methodology and workflow process for providing translational research informatics support," *Journal of Biomedical Informatics*, vol. 42, no. 2, pp. 377–381, 2009.
- [28] "Promis® instrument development and validation scientific standards version 2.0," 2013, revised May 2013.
- [29] E. Finell and J. Nätti, "The combined effect of poor perceived indoor environmental quality and psychosocial stressors on long-term sickness absence in the workplace: A follow-up study," *International journal of environmental research and public health*, vol. 16, no. 24, p. 4997, 2019.
- [30] T. Tuuminen, M. Andersson, S. Hyvönen, J. Lohi, and K. Vaali, "Indoor air nontoxicity should be proven with special techniques prior claiming that it may cause a variety of mental disorders," *International Journal of Hygiene and Environmental Health*, vol. 229, pp. 113 545–113 545, 2020.
- [31] "Nbr 16401-3: instalações de ar-condicionado, sistemas centrais e unitários: parte 3: qualidade do ar interior," 2008.
- [32] Y.-S. Park and S. Lek, "Artificial neural networks: Multilayer perceptron for ecological modeling," in *Developments in environmental modelling*. Elsevier, 2016, vol. 28, pp. 123–140.
- [33] S. S. Haykin *et al.*, "Neural networks and learning machines." 2009.
- [34] N. L. Costa, M. D. Lima, and R. Barbosa, "Evaluation of feature selection methods based on artificial neural network weights," *Expert Systems with Applications*, vol. 168, p. 114312, 2021.
- [35] J. D. Olden and D. A. Jackson, "Illuminating the "black box": a randomization approach for understanding variable contributions in artificial neural networks," *Ecological modelling*, vol. 154, no. 1-2, pp. 135–150, 2002.
- [36] H. Hewamalage, C. Bergmeir, and K. Bandara, "Recurrent neural networks for time series forecasting: Current status and future directions," *International Journal of Forecasting*, vol. 37, no. 1, pp. 388–427, 2021.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [38] M. Grządzielewska, "Using machine learning in burnout prediction: A survey," *Child and Adolescent Social Work Journal*, vol. 38, no. 2, pp. 175–180, 2021.

Trust Management Framework for Multi-Robot Systems

Daniel Vojnar, Adela Bierska, and Barbora Buhnova

ORCID: 0009-0009-1053-208X, 0009-0009-3119-530X, 0000-0003-4205-101X

Masaryk University, Faculty of Informatics

Brno, Czech Republic

Email: {vojnar, bierska, buhnova}@mail.muni.cz

Abstract—As autonomous technologies blend into a variety of industries, the employment of Multi-Robot Systems (MRS) in complex tasks is becoming increasingly prevalent. These systems, characterized by their distributed intelligence and collaborative capabilities, are now largely deployed in executing critical missions ranging from environmental exploration to disaster response. To support the trustworthy execution of such collaborative multi-robot missions, the existence of an underlying trust management framework is becoming imperative, given the rising risks of malicious intruders attempting to join the MRS.

This paper proposes a Trust Management Framework (TMF) for Multi-Robot Systems, to bridge the gap of scarce trust management support for MRS. By analyzing trust dynamics and integrating direct and indirect trust with contextual data, we address the trust vulnerabilities inherent to MRS. Our primary contribution is the development of a novel TMF significantly supporting the trust, security and reliability of MRS. The paper systematically outlines the evolution of our TMF, from theoretical underpinnings to an integrated solution, and the discussion of its impact on the trustworthiness of cooperative robotic networks.

I. INTRODUCTION

THE proliferation of intelligent and independent devices, vehicles, and robots is reshaping our everyday lives [1], [2]. This advancement leads to fully digital environments that can collaborate, compete, or even pose a threat [3], [4], [5]. Networks of multiple robots are advancing to take on human responsibilities in everyday activities and hazardous missions, including exploring underwater terrains, responding to natural disasters, and carrying out military operations [2].

The strength of Multi-Robot Systems (MRS) [6] lies in their distributed nature. Robots with unique capabilities possess different strengths and limitations, which enhances task success and system robustness. Robots have multiple sensors to sense their surroundings and plan the steps based on the mission and actual data. Communication is vital in these systems, enabling them to collaborate and adapt their behavior based on agreement and adaptive decision-making, which at the same time makes these systems vulnerable to robot misbehavior, whether unintentional or underpinned by hidden malicious intents.

Despite the increasing digitalization and interconnectedness of MRS, the integration of Trust Management Frameworks (TMF) [7] tailored for these systems remains largely unexplored. This gap exposes MRS to vulnerabilities, risking their robustness and resilience against attacks. Current trust

management approaches, while foundational, fall short in MRS contexts and missing understanding of their applicability and limitations in the context of MRS. These approaches also need to pay more attention to prioritizing individual safety of the members within these systems and the surrounding ecosystems.

To bridge this gap, this paper proposes a tailored Trust Management Framework (TMF) for Multi-Robot Systems (MRS), with the intent to overcome the discussed limitations by proposing a TMF that enhances the trustworthiness and safety of MRS through comprehensive trust inputs, including reputation and peer opinion, thereby ensuring the integrity and resilience of MRS in complex digital environments. To this end, we base our design on a general baseline TMF [7], to which we map the mechanisms necessary to prevent trust attacks in MRS scenarios, as collected in our previous work [8]. Specifically, we focus on the mechanisms necessary to support direct and indirect trust while paying attention to the contextual information necessary in trust decisions. Our main contribution is the creation of a novel Trust Management Framework to prevent vulnerabilities and enhance the safety of multi-robot systems. The proposed TMF is intended for feasible deployment and utilization while ensuring optimal trust and safety measures.

The rest of the paper is structured as follows. After reviewing related work in Section II, the background definition of multi-robot systems and robot capabilities for our specific purpose are provided in Section III. Section IV presents the methodology of the TMF design process, followed with individual components of the TMF design in Section V and VI. The resulting TMF for MRS is presented in Section VII, together with the discussion in Section VIII and conclusion in Section IX.

II. RELATED WORK

Machine-to-machine trust is crucial for successfully implementing fully autonomous systems involving multi-robot systems. The practical implementation of such trust was initially explored by Yang and Parasuraman in 2021 [9], who developed an agent trust model to facilitate cooperation among heterogeneous multi-robots. Their research focuses on trust evaluation based on the robot's needs (safety, basic, teaming, and capability). However, it could be enhanced by

other aspects like reputation or peer opinion, which could help provide solutions for unaddressed system vulnerabilities.

Recent studies delve deeper into the theoretical foundations of trust in autonomous systems [10], which focus on robot actions' reliability, predictability, and transparency and allow systems to form relationships. Besides, many studies can be found in the SIoT (Social Internet of Things) [11], which mimic human relationships. The study [12] focuses on applying aerial drone swarms to improve public safety by crowd monitoring and disaster response [13], requiring robust and reliable social interactions among robots. The [14] focused on trust-based mechanisms based on the blockchain model. Although their solution has outperformed various measures such as throughput, latency, accuracy, and block updating limits, it is based only on the local trust of its neighbors. Transitive relationships could improve the algorithm, degrading its performance but providing a more secure approach.

In contrast, the study [15] addresses the usage of Hidden Markov Model (HMM) in trust management for underwater robotics. This model enables the measurement of the trustworthiness of the sensor nodes to combat malicious or internal attacks. In this direction again, peer opinion could extend the robustness of trust management based on the HMM approach.

Ground-based networks, such as Vehicular Ad-hoc Networks (VANETs), allow vehicles to form human-like relationships, improving navigational aids and emergency response capabilities and ensuring safety for drivers and pedestrians [16]. Aslan and Sen designed a dynamic trust management model for vehicular ad-hoc networks [17] where the trustworthiness of vehicles is assessed using the vehicle trust value based on the data trust values of their event messages. This helps to establish a more reliable trust management framework with a combined trust model. While the model is a good source for inspiration, given its different context, it falls short in supporting details necessary in trust management in collaborative robotic missions (integrating the details of the mission, robot role in it, directives of the leader, etc.).

III. BACKGROUND

This section introduces the reader to multi-robot systems' essential terminology and background.

A. Multi-Robot Systems

The effectiveness of multi-robot systems is attributable to their composition of diverse robots, each enriched with distinct capabilities, enabling a synergistic collaboration that surpasses the abilities of individual robots. This significantly improves the success of a task and relaxes the dependence on a single robot, focusing on the cooperation of many distinct robots instead. As a result, the system becomes more robust and less likely to fail if one of the robots experiences an issue [18]. The distribution of robot types in the system is also essential, as different types of robots have different capabilities and limitations. Communication is a crucial aspect of multi-robot systems, which can be either seamless or introduce

challenges, depending on the compatibility of robot communication mechanisms and protocols. Yet, technology exists to facilitate communication between different communication protocols [19]. Efficient communication can significantly enhance the system's capabilities and increase its efficiency. Additionally, robots in multi-robot systems can recognize each other, known as kin recognition [20], and understand peer capabilities and limitations, further contributing to successful task completion.

B. Characteristics of a Robot in MRS

Inspired by [21], robot characteristics can be summarized as follows.

a) Capabilities: Robots are aware of their surroundings. They can generate and store information and combine it in real-time with new knowledge. They can plan their steps based on the mission and actual data. They have a rechargeable battery with limited capacity. They can move, orientate in the environment, and generate a trajectory considering detected obstacles. They are often capable of grabbing and carrying an (appropriately sized) item.

b) Sensors: Robots can sense their position using GPS and IMU. They can have any combination of additional sensors, such as LIDAR, radar, camera, microphone, temperature sensors, or range finders [22]. They can detect and diagnose internal faults and battery level.

c) Communication: Robots can communicate locally with other robots if the communication method is compatible. They can communicate on a private, ciphered channel with other robots with access to it. If possible, they can connect to a shared source of information and exchange data with it, both ways (download and upload). They can also store or backup data on remote storage. However, these remote connections may be impossible due to the bad conditions of the surrounding environment, inaccessible channels, or being too far away.

d) Collaboration: The robots are capable of collaboration with another robot or group of robots. They can consider the actions of other individuals while planning. They can participate in decision-making and adapt their behavior to the agreement. They can process and store information obtained from other robots.

IV. METHODOLOGY

The trust management concept is well-defined in various domains such as Internet of Everything or Social Internet of Things [7] but remains unexplored in multi-robot systems. With the ever-increasing digitalization and communication between systems, there is a need to ensure integration of appropriate trust management into such systems to increase their robustness and resilience for their safer integration in our digitalized society. To obtain an applicable solution for trust management in MRS, our Trust-Management Framework (TMF) design process is based on an existing trust framework, which is scrutinized from two directions to (1) validate whether each of its current components is necessary for trust

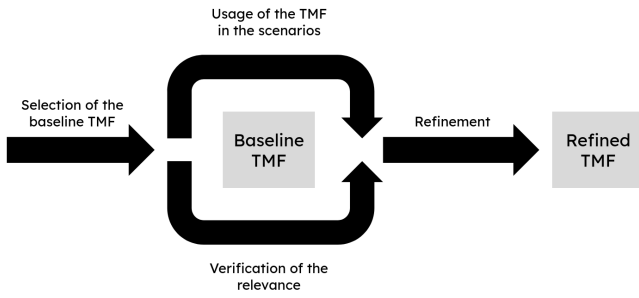


Fig. 1. Schema of the methodology.

management in MRS, and (2) identify missing properties and mechanisms specific to the context of MRS. Based on this analysis, a TMF for MRS solution is designed. These steps are visualized in Figure 1 and described below.

1) *Select the baseline TMF*: First, we identified a baseline trust management framework, with the intention to lay out a holistic foundation that can be then filtered and extended based on the specifics of MRS. To this end, we have reached for a more general context of the Internet of Everything (and Social Internet of Things) and chose a TMF designed as a collection of concepts from a broad literature review [7], [23]. In the TMF, depicted in Figure 2, trust of one agent (trustor) in another (trustee) is being built via combining the mechanisms of direct trust (based on local and highly context-specific experience with the other agent) and indirect trust (possibly global reputation of the agent) [24]. The two are later combined via the component of trust decision. The TMF components can be described as follows [7].

- *Direct Trust* constitutes an individual assessment by the trustor, derived from direct interaction with the trustee. It specifically stems from a combination of present and past experience (direct observation) that the trustor has with the trustee. Consequently, it is imperative to establish mechanisms that evaluate the experience during real-time interactions between the trustor and the trustee (potentially emulating human cognitive functions) without exposing its vulnerabilities.
- *Indirect Trust* is deduced from propagated opinions across various trust paths to assess a trustee's reputation. The principal sources are the trustor's trusted peers (their opinions and recommendations) and an overseeing authority responsible for monitoring the trustee's reputation within the network. Accordingly, there must be systems to update the reputation and propagate it throughout the ecosystem.
- *Context Information* shapes the trust decision to mirror the trustor's present circumstances during the decision-making process (for instance, the trustor's vulnerability during interactions with the trustee, the risks involved, the accountability of the trustee in the event of malicious actions, and whether any potential damage is reversible or subject to compensation, ensuring reparability). Fur-

thermore, broader contextual information also influences the direct and indirect trust computation.

2) *Analysis of the baseline TMF in the context of MRS*: To detect the necessary areas for change of the baseline TMF, we investigated its coverage of MRS context, which was done in two steps:

- *Relevance of existing components*. In the first step, we verified that the TMF contains no parts that could be irrelevant to the MRS. For each component in the TMF, we asked: *How can this component (direct trust/indirect trust/context information) be used to promote trust management in MRS?* and compiled answers with the help of knowledge of MRS from existing literature. A sufficient answer to this question denotes that the discussed component has its yield and, therefore, should remain in some form in the final solution.
- *Missing concepts and blind spots*. The second part of the analysis consists of detecting gaps in the coverage of the MRS context, which could cause the TMS to be insufficient. We went through the collection of the scenarios and vulnerabilities in multi-robot systems and its application in swarm robotics [8], and for each scenario and component, we asked: *How can this component be used to prevent or mitigate this type of trust attack in MRS? Is there anything missing in the original TMF that would ease this process?* Exhaustive answers to these questions serve as a base for updating the baseline TMF in the next step.

3) *Refined TMF for MRS Design*: Ultimately, we use the acquired knowledge to refine the baseline TMF and define the resulting TMF for MRS. The components of the baseline TMF, which were verified as valuable, serve as the base for the refined TMF. From the scenario analysis, we collect observations of the necessary robot support in trustworthy mission execution and cluster them by their source or subject. Each of the clusters gets assigned a component in the original TMF that it shall extend, or an approval to form a new component in the resulting TMF. Using these clusters, we refine the structure of the assigned components to cover them sufficiently.

V. ANALYSIS OF THE BASELINE TMF

This section examines the baseline Trust-Management Framework (TMF), as depicted in Figure 2, to validate whether each of its current components has its role in the context of trust management in Multi-Robot Systems (MRS).

A. Direct Trust

Direct trust is the main component by which the robot can judge other robots. Direct trust consists of two parts: present experience [25] and past experience [26]. Robots can use their current experience in MRS to evaluate their direct interactions with other robots. This can help to give more weight to the ongoing interaction. If a robot only bases its decisions on past interactions with another robot (recorded by past experience)

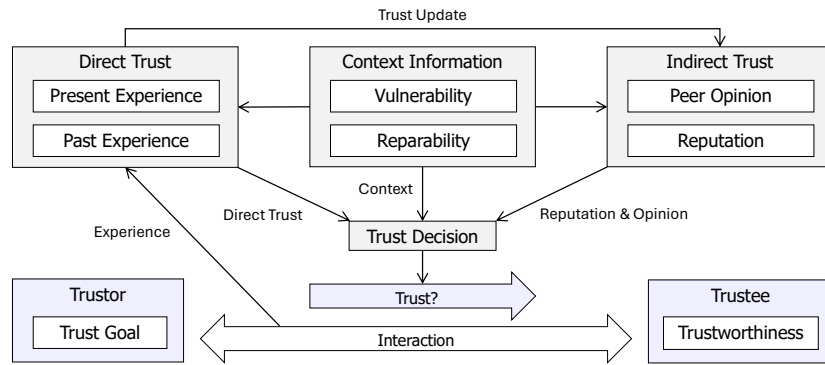


Fig. 2. Baseline Trust Management Framework [7].

and the other robot misbehaves in the current situation, the system might not recognize the misbehavior, and an error could occur. The trap of the experience mentioned above is a crucial component in MRS, as there are many one-to-one interactions between robots, which can provide the robot with significant experience that can help it evaluate trust. This shows that both direct trust components of the baseline TMF are relevant and contribute to correct functioning.

B. Indirect Trust

In the MRS, robots can use direct and indirect trust [27] to evaluate each other. While direct trust comes from personal interaction, indirect trust comes from reports, recommendations, and reputation mediated by others. Robots can leverage indirect trust by relying on peer opinions [28] to decide whether to trust an individual. They can also establish a hierarchy based on individuals' reputations [28], which can increase when a robot performs its tasks well. If a robot completes most of a task, even if not explicitly designed for it, its reputation can reflect it and increase. All baseline TMF aspects are relevant and suitable for the MRS.

C. Context Information

For MRS, context information can carry the same significance as other systems. It can serve as an essential but non-negligible aspect in the process of making trustworthy decisions. Context information can aid in evaluating the suitability, risks [29], and potential advantages of an action or interaction within a specific environment [30] and with specific objects. Environmental factors, information sensitivity [31], and the likelihood of recovering [7] from a risky action can be considered in this context. Additionally, the types of robots in our environment can influence MRS, as each individual may be suited to different tasks, leading to varying levels of security perception. In summary, all the components of the context information baseline TMF remain relevant also in the context of MRS.

VI. TRUST SCENARIOS MAPPING TO THE BASELINE TMF

This section presents the analysis of the baseline TMF with the intention to identify missing concepts and blind spots

relevant to the context of MRS. It focuses on individual vulnerabilities occurring in multi-robot systems, as collected in [8]. The text is divided into several sections, each explaining the issue and providing a detailed description of the trust-management solution and its mapping to the baseline TMF.

A. Information Manipulation or Ignoring

The improper modification and dissemination of information can cause significant problems during a mission. The transmitted data can be divided into three categories of actions related to erroneous information: individual, external, and passing information [8].

The presence of a Trust Management Framework (TMF) can support the trustor to cross-check information obtained from the trustee with other sources. Therefore, it can point out inconsistencies in the information, which can lead to discarding false information and thereby foiling the attack. The type of mission sets the sensitivity of the data, so it should be considered in the context part of the TMF for MRS. Also, only insiders (trusted peers) may recognize some defects in information as they have the necessary permissions to access, so in the TMF for MRS, they should be differentiated from random passersbys.

1) *Individual Information:* Let us first examine the TMF robustness against the attack of a robot sharing false individual information. The purpose of this attack is for the robot to share false information about itself that others cannot verify, such as its battery level. The TMF can be very beneficial in this case, as it can use its recorded past interactions with the robot, compare information with other sources [32], and send a test message to determine if the robot is generally truthful. If the system is centralized [33], the robot may have to be asked to answer to a higher authority with more information about itself.

2) *External Information:* Information manipulation or ignoring attacks on external information are similar to the previous type of attack. In this scenario, the robot observes [34] its surroundings and shares false information about other robots or objects in the environment, which can however be verifiable by others. Since the information is verifiable, other robots can verify the information and update the source robot

trust scores accordingly. The impact on behavior depends on environmental conditions and possibly on its changes (e.g., weather, moving objects), which are currently not included in the baseline TMF.

3) *Passing the Information*: In this type of attack, the robot receives information from other sources, such as another robot or a server, and then passes it along to the next recipient, which is the moment when it can change it to false information. In this scenario, the robot is not the original source of the information but rather a mediator that can modify the transmitted data. An important aspect of robot interactions is the direct experience with a robot that consistently and accurately transmits information without making unwarranted modifications. Trust values from other peers can be received through feedback [32], reputation ratings [35], or recommendations [36]. The authority can select the primary message distributors based on their trust level if the message is not broadcast. When the robot transmits sensitive data, like navigational information, the communication should be more reliable, which can have a different impact on the trust score [31].

B. Manipulation with Communication Channels

Effective communication and information exchange are crucial for mission success. However, vulnerabilities in data passing and storage can be exploited. Even a minor leak of swarm position or resources may lead to significant losses. Insider attacks and changes made to shared information can be significant risks. Also, interference with communication and signals could threaten the mission [8].

Observation of malicious behavior against communication channels by the trustor or their peers, followed by the isolation of the trustee, may stop those attacks from spreading. However, detection of this behavior gets more complex with various permissions and chosen communication channels. Therefore, the TMF for MRS should go into more detail with the context of the mission and roles of the robots.

1) *Leak of Information*: The trustor needs to observe trustee's history of managing and transmitting sensitive information without exposing it. To detect vulnerable behaviors with the information, storing every robot's action globally and locally is necessary. It will be a marked comparison of these historical actions [37] to see if they are the same or if there are any incorrect things.

Indirect trust can enhance information's credibility by allowing other robots' feedback and reports. This would also enable the implementation of an external security system to monitor unusual communication patterns [37]. Robots involved in previous security breaches or have indirect evidence of such behavior would be considered less trustworthy, even if there is no direct observation of their actions.

The sensitivity of the transferred information is essential for evaluating the manipulation with it as the more sensitive data should be protected better. In the current TMF, this is not completely covered as the context part aims more at the trustor itself than its mission.

2) *Changing of Shared Information*: Many algorithms often work together using shared information. However, this shared source can be vulnerable to manipulation, leading the group to work with false information.

The trustor should observe the trustee's correct manipulation in updating and handling the information, which enhances the group's problem-solving capabilities. If the trustee is caught making unauthorized or harmful changes in the past, it should take shape in the future evaluations of their trustworthiness.

The shared information should be checked and compared with previous versions, and the changes that were caused to the data should be evaluated. If other robots suspect a change made to the information, they can share that with each other.

3) *Restrain Access*: The work of MRS relies on distributed communication and information sharing, making interference a real threat. Malicious robots can disrupt signal transmission by overloading the communication channel with irrelevant messages.

There are always multiple communication channels between robots; the higher the trust score of a robot, the more channels are available. The higher the level of the communication channel, the more messages can pass through it with a higher level of confidentiality. The origin and content of messages should be observed to detect overflowing by irrelevant or repeating messages. The evaluation is based on the context of the mission and its presumed communication.

Malicious robots can also disturb communication by noise or physical destruction of transportation medium. Their behavior and sent signals should be observed and evaluated.

The trustor may verify sent messages with peers on other more secure channels to which only they can access.

C. False Performance Promises

In time-critical missions, there is a risk that some individuals may intentionally fail to complete tasks on time. Other robots failing to identify and address this malicious behavior can lead to unnecessary time loss and mission failure. The more time-critical the mission, the more vulnerable it is to this attack. The suitability of the environment and the trustee for the mission can also impact the time limit that the trustor is willing to wait for the mission to be completed, which is connected to the vulnerability of the whole system and the individual waiting for the completion of an action to continue with his task. The time and status of the mission are reflected in the timeout for action. These aspects should be included in the baseline Trust Management Framework (TMF).

The robot's past experience determines if a delay is typical. This past experience can be shared among all system robots, enabling them to share opinions about the targeted individual.

If robots are visibly damaged or noticeably delayed in their work, it has to be factored into the trust decision. Furthermore, the delay can be compared to a pre-estimated time limit.

D. Authority Misusage

In swarm systems with hierarchical leadership, malicious robots in leader roles can manipulate others to do unwanted

things, making detection and elimination harder. Malicious leaders may change permissions and force subordinates into unwanted actions, endangering the mission.

This scenario is the least covered by the baseline TMF as it does not recognize the trustee's relation to the mission and the role of the leader as a whole. But it can still detect suspicious behavior of the leader and cause their demotion.

1) *Permission Regulations*: Each robot should have permission to access resources set according to its role or mission. A malicious robot with authority may change the rights of other robots, denying access to information or resources. If these permissions are not granted or are taken away without clear reason, it should be considered suspicious. These rights assignments could be taken as direct experience within the TMF.

To evaluate their own rights, the trustor should take into account the value and sensitivity of the data used within the mission. It needs to be as objective as possible because even if the robot thinks it needs a particular resource, it does not mean there is no way to solve the mission without it. The baseline TMF covers this sensitivity to some extent through context information.

When introducing the leader hierarchy, the trustor can ask sub or superior leaders for their opinion. Leaders' opinions may have a higher impact on the decision or even be unquestionable. Therefore, this weight could be somehow recognized in the TMF.

Also, the trustor can observe inadequate changes in the permissions of their peers. In that case, it may signify the leader's (trustee's) incompetency or malicious intentions, which are both problematic for the mission's success. The baseline TMF covers this by the peer opinion part.

2) *Task Allocation*: Getting orders from the leader is usually done via direct communication, and if the trustor knows and understands the mission's aims and plans, they can directly evaluate this kind of interaction. The orders should be consistent with previous ones and any changes should have a reason. The leader should have a plan and reasons for task assignment. If they encourage the trustor to do irrelevant or dangerous tasks, they should provide more details, eventually even evidence that a particular approach does not lead to the destruction of the society.

As the missions of MRSs are, by definition, collaborative, so orders from leaders should take into account the plans of other group members and complement them efficiently. Therefore, references from peers and leaders within the same mission are more valuable for detecting this attack. If the assigned task is irrelevant or even complicates or devalues the work of other group members, it should be considered suspicious. However, similarly to the previous section, the task evaluation should be as objective as possible because the mission may be in a critical state, which requires more extreme solutions.

E. Physical Attacks

In robotics, physical attacks can range from inter-robot communication manipulation to the destruction of robots or

even the kidnapping of individual units. These attacks can significantly delay a robot's mission, including stealthy tactics or leading robots into traps. Additionally, changing or destroying the environment can make it difficult for robots to navigate and complete their mission.

Components in the TMF cover all necessary information to detect possible physical attacks. The trustor evaluates their vulnerability and reparability and, based on past or present destructive behavior of the trustee or their bad reputation, may decide to avoid them or even isolate them.

1) *Robot Destruction*: Aggressivity and destructive intentions may be long-term, especially in the case of trustee malfunction. Therefore, the trustor should consider previous endangering behavior covered by past experience in the TMF. Present experience is also important, the trustor should observe the possible physical superiority of the trustee and their behavior. Even false information or risky instructions obtained from the trustee may mean a deliberate effort to destroy the trustor.

Physical attacks are usually critical and unmistakable with other actions. Any peer can similarly evaluate the situation and refer it to the trustor. Therefore, the opinion of robots with and without knowledge of the mission and context can be considered valuable without much doubt.

The actual state and capabilities of the trustor directly influence its ability to defend itself against the destructive trustee. Also, the trustor's importance for the mission should be considered in the final trust decision. The surrounding environment and mission type influence the view of the trustee's communication. For example, risky instruction may be more tolerable in dangerous environments under time pressure.

2) *Kidnapping/Capture*: Similarly to the destruction attacks, kidnapping may be predicted by the uneven physical capabilities of the trustor and trustee and by suspicious communication. Exhortation to move to unknown places should especially be critically evaluated.

The risk of being kidnapped impels the trustor to consider their value. They may carry secret information or know-how that may be misused in the wrong hands. Otherwise, the context may be weighted as in case of possible destruction.

The escaped peers will provide information about the kidnapper's behavior. Therefore, the trustor should consider the peer's opinion. Also, the size and strength of the surrounding group can be considered, as the trustor alone is more vulnerable to kidnapping than the whole group.

3) *Changing or Destruction of the Environment*: Spotting trustee to manipulate their surroundings inadequately may signal their destructive intentions. Also, observing new environmental changes may lead to suspecting nearby robots.

In addition to warnings from witnesses of the malicious behavior, the trustor's peers may also provide information about the environment's last state and the items' positions. Based on this, any changes may be observed without the previous visit to the particular place.

The environment may contain items of various value, and the trustor should evaluate their importance according to the

TABLE I
COLLECTED OBSERVATIONS SORTED BY TMF PART

Direct Trust	Context Information	Indirect Trust
Trustee's motion	State of the environment	Communication with other agents
Communication with trustee	Trustor's experience with similar environment	Feedback loops
Trustor's historical experience with trustee	Static vs dynamic object	Reputation
Accuracy of provided information	Sensitivity of data	Involvement in previous security breaches
Integrity of provided information	Riskiness of the mission	History of shared information source
Trustee's behavior	Discovery of information leak	Other leaders' opinion
Trustee's knowledge	Credibility of provided information	Missing peers
Unauthorized or harmful change of shared information	Trustees's access to information	Leader's opinion
Amount of sent information from trustee	Relevance of information	Destroyed peers
Capabilities of trustee	Author of information	
State of trustee	Time left	
Trustees opinions	Trustor's access to information	
	Threat to other robots	
	Efficiency	
	Mission's state	
	Trustor's state	
	Trustor's capabilities	
	Changes of environment	

mission. Changing or destruction of the environment may not always mean bad intentions. There could be, for example, some obstacles or dangerous objects intended to be displaced or destroyed.

F. Attacks on Internal Intelligence

Machine learning models have vulnerabilities that lead to biased or irrational behavior. To prevent this, it is important to stop false data from reaching the robot's learning model and to verify information with trusted sources. It is also crucial to monitor the robot's behavior for signs of exploitation of known vulnerabilities in its control mechanisms [38]. The robots have past experiences stored in their memory, which includes their encounters with malicious data. The distributed intelligence of the entire group, including the robots' internal intelligence, must be well-protected to prevent attacks by malicious entities. This involves addressing complex decision-making under uncertainty and minimizing information leakage [39], [40].

G. Decision Making Attacks

The decision-making process of robots can be influenced by various types of attacks, such as contrarian behavior, wishy-washy attitudes, following a sect [41], and going along with the majority opinion [8]. It is essential to monitor the trustee's opinions over time to detect any potentially harmful patterns aimed at manipulating decisions. Opinions can be affected by the current context, the mission status, and the trustee's relationship to it, which are essential factors currently not considered. Peer opinions can also carry weight if the peer is trusted and working with the same information and mission. Therefore, understanding the trustor-peer relationship can enhance the effectiveness of the decision-making process. Bots can omit or exclude a trustee from voting if past experiences show the trustee attempted to abstain. The history of individuals plays a significant role in preventing these kinds of attacks. When evaluating opinions, it is crucial to consider the context and recognize that initially controversial opinions may be supported by concrete steps or facts.

H. Summary of Observations

Overall, we have identified 39 observations that characterize the specifics of MRS context in trust management across all the examined scenarios. The collected observations were then clustered by their affiliation to a particular component of the baseline TMF, as presented in Table I.

Main topics missing by baseline TMF are mission information (e.g. its state, plan, kind of data) and different levels inside of the mission hierarchy (leaders, team members, other passerby).

VII. TRUST MANAGEMENT FRAMEWORK FOR MULTI-ROBOT SYSTEMS

This section presents the proposed Trust Management Framework (TMF) for Multi-Robot Systems (MRS), reflecting the observations from Section V and VI, as well as the mapping of the observations to TMF components as presented in Table I. The resulting TMF for MRS is presented in Figure 3.

In summary, the observations led to a more detailed breakdown of the *direct trust* TMF component, where the *present experience* is divided into *communication*, the accuracy and credibility of which should be evaluated by the *trustor* as the *trustee* may easily lie, and *observation*, that does not have to be directly questioned because it comes from a reliable source (*trustor* itself).

Mission information expanded the *context information* component, as the state of the mission directly influences the need for collaboration with other robots and overall safety requirements. The *mission information* consists of three main aspects: *mission state*, *time left*, and *data sensitivity*, described in detail below.

Besides, as the robots may play different roles within the mission with different roles to the *trustor*, the TMF was extended to reflect these roles. The rest of this section presents the TMF components in a structured way.

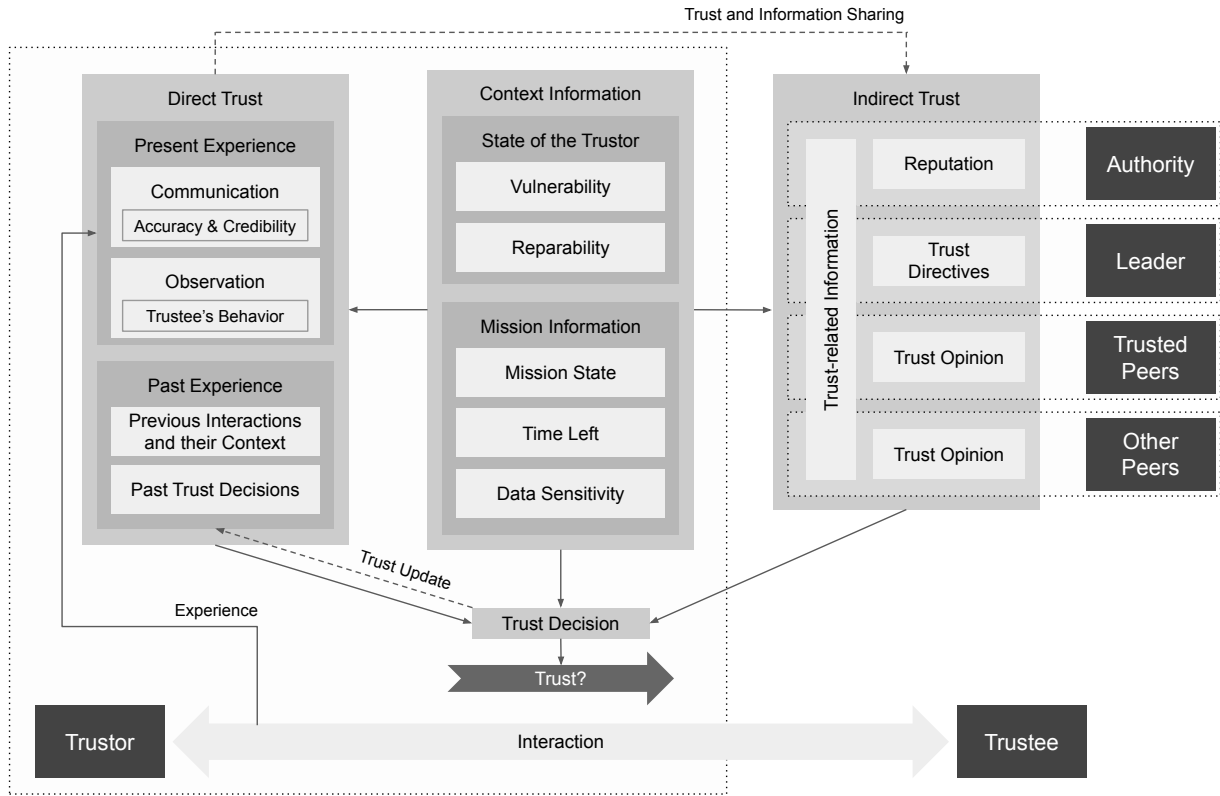


Fig. 3. Trust Management Framework for Multi-Robot Systems.

A. Actors

1) *Trustor*: represents the robot in question who is in the process of making a decision whether to trust another robot (*trustee*). The TMF is in place to support this *trust decision*.

2) *Trustee*: represents a robot (typically outside the scope of control by the *trustor*) that is interacting in some way with the *trustor* who has limited certainty about its trustworthiness and might feel vulnerable in its presence.

3) *Authority*: sits on top of the trust hierarchy and might take form of a set of regulations and directives (does not need to be represented as a robot itself).

4) *Leader*: is superior to the *trustor* and controls the mission the *trustor* is deployed to complete. In this regard, its trust directives should have higher weight than peer opinion or even be, by definition, unquestionable.

5) *Trusted Peers*: represent the robots whose trust opinion has higher weight in *trust decision* than that of other robots in the system, as they form a peer (friendship) group with the *trustor*, whose trust they have previously gained.

6) *Other Peers*: represent other robots who might be relevant in the context of the mission that the *trustor* is aiming to complete but who do not benefit from pre-gained trust of the *trustor*.

B. Direct Trust

The direct trust component consists of the *present* and *past experience* of the *trustor* with the *trustee*, being fed by their ongoing interaction. The *present experience* is influenced by the *communicated* information (its accuracy and credibility) and trustor's direct *observation* of trustee actions and behavior. The *past experience* reflects the experience from *previous interactions*, together with their *context*, such as the role of the trustee in the mission that the past experience relates to (i.e., how dependent the trustor was on the trustee at that point in time to complete its mission). Besides, the *past experience* also records *past trust decisions* made by the *trustor* about the *trustee*.

C. Indirect Trust

The indirect trust component is specific to the role it originates from. While the *authority* governs the global *reputation* of each MRS member, the *leader* gives *trust directives* (e.g., the necessity to trust a certain robot), the peers (both the *trusted peers* and *other peers*) only share their trust opinion, together with supplementary *trust-related information*, which can be shared by any of these actors. The *trust-related information* can, for instance, consist of the elements of peers' subjective *past experience* with the *trustee* in question.

D. Context Information

The *context information* characterizes the context of the *trustor* that influences its *trust decision*. This includes its *vulnerability* and *reparability* as in the baseline TMF, but it also includes *information about the mission* that the *trustor* is deployed to complete (namely the *mission state*, *time left*, and *data sensitivity*). *Mission state* stands for mission progress so far (e.g., environment explored, part of the resources lost) and its future plan (e.g., collecting environmental information, returning to the base). *Time left* records how much time is left to complete the mission. The *mission state* and *time left* directly influence the robot's behavior, decision-making process, the need for collaboration with other robots, and overall security and safety requirements. *Data sensitivity* mainly impacts data transmission and communication—the higher the sensitivity level, the higher trust score of a *trustee/peer* is necessary to approve data communication.

E. Trust Decision

The *trust decision* is made on the basis of the *direct trust* consideration, *indirect trust* inputs, and the current state of the *context* in the particular moment when the *trust decision* is to be made. The result is stored in the *past experience* part of the *direct trust* component (as it is subjective to the *trustor*) and propagated to the rest of the network via *trust and information sharing*.

VIII. DISCUSSION

With the advent of Industry 4.0, Smart Cities, and other advanced contexts [42], [43], MRS are becoming increasingly present in everyday life, where increasingly more tasks rely on them. Therefore, their trustworthiness and safety needs to be taken more seriously. Our work brings ideas on preventing attacks even before they occur by carefully selecting trusted system members. By avoiding suspicious individuals, the whole MRS can become more trustworthy, and hence, it can participate in more critical tasks, including interaction with vulnerable human beings.

The selection of trusted peers cannot be done imprudently as the MRS missions are intended to be done by collaboration between multiple robots. An extensive reduction in the pool size of trusted peers might thus lead to lower efficiency or even failure of the mission.

The proposed TMF for MRS brings structured information and methods to make trust decisions consciously by considering relevant data and the real-time state of the mission and the environment. As the amount of this data may be overwhelming, its processing should be considered in the early phase of the system design. The TMF should be narrowly targeted to MRS to cover all their needs, but it should also be flexible enough to adapt to any of the wide range of MRS types.

A. Threats to Validity

To minimize the threats to validity, the proposed TMF for MRS is built on two main sources, covered by careful

methodological design to minimize the risks of compromising its quality. To this end, we have opted for an incremental design, starting from an existing baseline TMF and scrutinized its components in both the direction of the necessity of existing components and their sufficiency to cover possible MRS scenarios.

B. Application of the Proposed Framework

The framework shows which facts should be considered when deciding trust between the trustor and the trustee. When designing a MRS, all of these items should be considered and covered to achieve the best result. Each part of the TMF should be adjusted to the context of the concrete robot, and the collected data should be structured to evaluate the trust effectively. The trustor needs to be assigned an initial trust score and relation to the newly encountered robot and robots from the initial group (considering its leader hierarchy). After that, trust should be evaluated regularly, even with robots deemed trustworthy, due to the trust score erosion over time (governed by the system authority). After each evaluation, the trustor should update records about the trustee in its memory and inform the rest of the system about its findings.

C. Future Directions

There is a multitude of research directions that can take the proposed TMF further, whether in terms of its application or extension. The first intended step on our side is to experiment with the framework in a variety of case studies, exploring the TMF in the context of heterogeneous MRS, where robots with vastly different designs, capabilities, and purposes must work together seamlessly. Next, we aim to explore how the TMF can be scaled to accommodate large and complex MRS, ensuring that trust management remains effective as the number of robots and interactions increases.

In the next phase, we intend to explore the role of the designed TMF in human-robot trust dynamics, especially in scenarios where human intervention or collaboration is necessary. This might include the need to address the ethical implications of trust decisions made by MRS and ensure compliance with emerging regulations in robotics and AI.

IX. CONCLUSION

In conclusion, this research contributes to the potential of Multi-Robot Systems (MRS) in our digital era, where their distributed intelligence and collaborative capabilities are being harnessed for a multitude of applications, from simple tasks to complex missions. The inherent strength of MRS, derived from the unique abilities and collective resilience of their robotic constituents, is limited by their susceptibility to misbehavior and security threats. Recognizing this challenge, this paper contributes to the solution by introducing a Trust Management Framework (TMF) for MRS designed to fortify the integrity and safety of MRS. Grounded in a baseline TMF, our proposed framework integrates MRS-related mechanisms fed by extensive research on MRS scenarios. The framework's emphasis on both direct and indirect trust, underpinned by

contextual awareness, represents a holistic solution towards safeguarding MRS against vulnerabilities and opening to door towards MSR integration in our interconnected world.

ACKNOWLEDGMENT

The work was supported by GAMU project "Forensic Support for Building Trust in Smart Software Ecosystems" (no. MUNI/G/1142/2022).

REFERENCES

- [1] R. Capilla, E. Cioroica, B. Buhnova, and J. Bosch, "On autonomous dynamic software ecosystems," *IEEE Transactions on Engineering Management*, vol. 69, no. 6, pp. 3633–3647, 2022. doi: 10.1109/TEM.2021.3116873
- [2] H. Bangui, B. Buhnova, D. Kusnirakova, and D. Halasz, "Trust management in social internet of things across domains," *Internet of Things*, vol. 23, p. 100833, 2023.
- [3] D. Halasz and B. Buhnova, "Rethinking safety in autonomous ecosystems." in *FedCSIS (Position Papers)*, 2022, pp. 81–87.
- [4] A. Bierska, B. Buhnova, and H. Bangui, "An integrated checklist for architecture design of critical software systems." in *FedCSIS (Position Papers)*, 2022, pp. 133–140.
- [5] M. Macac, S. Bojnak, and B. Buhnova, "Identification of unintentional perpetrator attack vectors using simulation games: A case study," in *2021 16th conference on computer science and intelligence systems (FedCSIS)*. IEEE, 2021, pp. 349–356.
- [6] Y. Rizk, M. Awad, and E. W. Tunstel, "Cooperative heterogeneous multi-robot systems: A survey," *ACM Computing Surveys (CSUR)*, vol. 52, no. 2, pp. 1–31, 2019. doi: 10.1145/3303848
- [7] B. Buhnova, "Trust management in the Internet of Everything," in *European Conference on Software Architecture. ECSA 2022 Tracks and Workshops*. Springer, 2023. doi: 10.1007/978-3-031-36889-9_10 pp. 123–137, preprint at <http://arxiv.org/abs/2212.14688>.
- [8] D. Vojnar, A. Bierska, and B. Buhnova, "Scenarios for trust management in swarm robotics," 2024.
- [9] Q. Yang and R. Parasuraman, "How can robots trust each other for better cooperation? a relative needs entropy based robot-robot trust assessment model," in *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2021. doi: 10.1109/SMC52423.2021.9659187 pp. 2656–2663.
- [10] P. Gangwani, A. Perez-Pons, and H. Upadhyay, "Evaluating trust management frameworks for wireless sensor networks," *Sensors*, vol. 24, no. 9, p. 2852, 2024. doi: 10.3390/s24092852
- [11] F. Amin and G. S. Choi, "Social pal: A combined platform for internet of things and social networks," in *2020 5th International Conference on Computer and Communication Systems (ICCCS)*. IEEE, 2020. doi: 10.1109/ICCCS49078.2020.9118579 pp. 786–790.
- [12] S. H. Alsamhi, O. Ma, M. S. Ansari, and S. K. Gupta, "Collaboration of drone and internet of public safety things in smart cities: An overview of qos and network performance optimization," *Drones*, vol. 3, no. 1, p. 13, 2019. doi: 10.3390/drones3010013
- [13] S. M. S. M. Daud, M. Y. P. M. Yusof, C. C. Heo, L. S. Khoo, M. K. C. Singh, M. S. Mahmood, and H. Nawawi, "Applications of drone in disaster management: A scoping review," *Science & Justice*, vol. 62, no. 1, pp. 30–42, 2022. doi: 10.1016/j.scijus.2021.11.002
- [14] G. Rathee, A. Kumar, C. A. Kerrache, and R. Iqbal, "A trust-based mechanism for drones in smart cities," *IET Smart Cities*, vol. 4, no. 4, pp. 255–264, 2022. doi: 10.1049/smc2.12039
- [15] M. M. Arifeen, D. Bhakta, S. R. H. Remu, M. M. Islam, M. Mahmud, and M. S. Kaiser, "Hidden markov model based trust management model for underwater wireless sensor networks," in *Proceedings Of The International Conference On Computing Advancements*, 2020. doi: 10.1145/3377049.3377054 pp. 1–5.
- [16] A. Deshpande, "Review of effective trust management systems in vanet environments," *International Journal of Grid and Distributed Computing*, vol. 14, no. 1, pp. 1771–1780, 2021.
- [17] M. Aslan and S. Sen, "A dynamic trust management model for vehicular ad hoc networks," *Vehicular Communications*, vol. 41, p. 100608, 2023. doi: 10.1016/j.vehcom.2023.100608
- [18] A. Gautam and S. Mohan, "A review of research in multi-robot systems," in *2012 IEEE 7th international conference on industrial and information systems (ICIIS)*. IEEE, 2012. doi: 10.1109/ICIInfS.2012.6304778 pp. 1–5.
- [19] F. Steele Jr and G. Thomas, "Directed stigmergy-based control for multi-robot systems," in *Proceedings of the ACM/IEEE international conference on Human-robot interaction*, 2007. doi: 10.1145/1228716.1228747 pp. 223–230.
- [20] K. Bolla, T. Kovacs, and G. Fazekas, "Compact image processing based kin recognition, distance measurement and identification method in a robot swarm," in *2010 International Joint Conference on Computational Cybernetics and Technical Informatics*. IEEE, 2010. doi: 10.1109/IC-CYB.2010.5491237 pp. 419–424.
- [21] O. Vermesan, R. Bahr, M. Ottella, M. Serrano, T. Karlsen, T. Wahlstrøm, H. E. Sand, M. Ashwathnarayan, and M. T. Gamba, "Internet of robotic things intelligent connectivity and platforms," *Frontiers in Robotics and AI*, vol. 7, p. 104, 2020. doi: 10.3389/frobt.2020.00104
- [22] A. Cowley, H.-C. Hsu, and C. J. Taylor, "Distributed sensor databases for multi-robot teams," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, vol. 1. IEEE, 2004. doi: 10.1109/ROBOT.2004.1307229 pp. 691–696.
- [23] S. Sagar, A. Mahmood, Q. Z. Sheng, J. K. Pabani, and W. E. Zhang, "Understanding the trustworthiness management in the social internet of things: A survey," *arXiv preprint arXiv:2202.03624*, 2022. doi: 10.1016/j.comnet.2024.110611
- [24] N. B. Truong, T.-W. Um, B. Zhou, and G. M. Lee, "From personal experience to global reputation for trust evaluation in the social internet of things," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, 2017. doi: 10.1109/GLOCOM.2017.8254523 pp. 1–7.
- [25] J. Lee and J. C. Oh, "A node-centric reputation computation algorithm on online social networks," *Applications of Social Media and Social Network Analysis*, pp. 1–22, 2015. doi: 10.1007/978-3-319-19003-7_1
- [26] C. Marche and M. Nitti, "Can we trust trust management systems?" *IoT*, vol. 3, no. 2, pp. 262–272, 2022. doi: 10.3390/iot3020015
- [27] R. Hussain, J. Lee, and S. Zeadally, "Trust in vanet: A survey of current solutions and future research opportunities," *IEEE transactions on intelligent transportation systems*, vol. 22, no. 5, pp. 2553–2571, 2020. doi: 10.1109/TITS.2020.2973715
- [28] B. Zhang, Z. Huang, and Y. Xiang, "A novel multiple-level trust management framework for wireless sensor networks," *Computer Networks*, vol. 72, pp. 45–61, 2014. doi: 10.1016/j.comnet.2014.06.015
- [29] H. El-Sayed, H. A. Ignatious, P. Kulkarni, and S. Bouktif, "Machine learning based trust management framework for vehicular networks," *Vehicular Communications*, vol. 25, p. 100256, 2020. doi: 10.1016/j.vehcom.2020.100256
- [30] S. A. Ghasempouri and B. T. Ladani, "Modeling trust and reputation systems in hostile environments," *Future Generation Computer Systems*, vol. 99, pp. 571–592, 2019. doi: 10.1016/j.future.2019.05.017
- [31] I. Ahmed, T. Nahar, S. S. Urmi, and K. A. Taher, "Protection of sensitive data in zero trust model," in *Proceedings of the international conference on computing advancements*, 2020. doi: 10.1145/3377049.3377114 pp. 1–5.
- [32] R. Govindaraj, P. Govindaraj, S. Chowdhury, D. Kim, D.-T. Tran, and A. N. Le, "A review on various applications of reputation based trust management," *International Journal of Interactive Mobile Technologies*, vol. 15, no. 10, 2021. doi: 10.3991/ijim.v15i10.21645
- [33] I. U. Din, K. A. Awan, A. Almgren, and B.-S. Kim, "Sharetrust: Centralized trust management mechanism for trustworthy resource sharing in industrial internet of things," *Computers and Electrical Engineering*, vol. 100, p. 108013, 2022. doi: 10.1016/j.compeleceng.2022.108013
- [34] F. Ishmanov, A. S. Malik, S. W. Kim, and B. Begalov, "Trust management system in wireless sensor networks: design considerations and research challenges," *Transactions on Emerging Telecommunications Technologies*, vol. 26, no. 2, pp. 107–130, 2015. doi: 10.1002/ett.2674
- [35] H. Yu, Z. Shen, C. Miao, C. Leung, and D. Niyato, "A survey of trust and reputation management systems in wireless communications," *Proceedings of the IEEE*, vol. 98, no. 10, pp. 1755–1772, 2010. doi: 10.1109/JPROC.2010.2059690
- [36] A. Abdul-Rahman and S. Hailes, "Supporting trust in virtual communities," in *Proceedings of the 33rd annual Hawaii international conference on system sciences*. IEEE, 2000. doi: 10.1109/HICSS.2000.926814 pp. 9–pp.

- [37] S. Ruohomaa and L. Kutvonen, "Trust management survey," in *International Conference on Trust Management*. Springer, 2005. doi: 10.1007/11429760_6 pp. 77–92.
- [38] C. Blum and D. Merkle, *Swarm intelligence: introduction and applications*. Springer Science & Business Media, 2008. ISBN 9783540740896, 3540740899
- [39] M. H. Nasir, S. A. Khan, M. M. Khan, and M. Fatima, "Swarm intelligence inspired intrusion detection systems—a systematic literature review," *Computer Networks*, vol. 205, p. 108708, 2022. doi: 10.1016/j.comnet.2021.108708
- [40] M. Srivatsa, S. Balfe, K. G. Paterson, and P. Rohatgi, "Trust management for secure information flows," in *Proceedings of the 15th ACM conference on Computer and communications security*, 2008. doi: 10.1145/1455770.1455794 pp. 175–188.
- [41] F. Canciani, M. S. Talamali, J. A. Marshall, T. Bose, and A. Reina, "Keep calm and vote on: Swarm resiliency in collective decision making," in *Proceedings of workshop resilient robot teams of the 2019 IEEE international conference on robotics and automation (ICRA 2019)*, vol. 4, 2019.
- [42] L. Wallezky, B. Buhnova, and L. Carrubbo, "Value-driven conceptualization of services in the smart city: a layered approach," *Social dynamics in a systems perspective*, pp. 85–98, 2018.
- [43] S. Chren, B. Rossi, B. Buhnova, and T. Pitner, "Reliability data for smart grids: Where the real data can be found," in *2018 smart city symposium prague (scsp)*. IEEE, 2018, pp. 1–6.

Enhancing Airbnb Price Predictions with Location-Based Data: A Case Study of Istanbul

Özgün Akalın
Computer Engineering
Galatasaray University
Istanbul, Turkey
0009-0001-2067-856X

Gülfem Isiklar Alptekin
Computer Engineering
Galatasaray University
Istanbul, Turkey
0000-0003-0146-1581

Abstract—Airbnb, a prominent online marketplace, facilitates short- and long-term rentals by connecting customers with property owners offering entire apartments or private rooms. Accurate price prediction is crucial for both the platform and rental property owners. Previous studies have primarily focused on statistical methods and pre-processing techniques, with limited exploration of the impact of location attributes. This paper aims to enhance price prediction models for Airbnb listings by incorporating location data. Utilizing data from InsideAirbnb for Istanbul, we implemented various data pre-processing techniques and enriched the dataset with location-specific information. Our findings show that incorporating these location-based features significantly improved model performance, increasing the adjusted R^2 metric by 22.5% and reducing Mean Absolute Error (MAE) by %10. This enhancement was achieved by using location-related index values and public transportation data provided by the Istanbul Metropolitan Municipality. These advancements can help property owners optimize rental prices and assist urban planners in making informed decisions about city infrastructure development.

Index Terms—Price prediction, regression, XGBoost, location features, Airbnb

I. INTRODUCTION

AIRBNB is an online platform where hosts rent their entire apartments or rooms to guests [1]. As of the end of 2023, it serves as a global marketplace with over 5 million hosts and 7.7 million active listings. Airbnb's offerings differ from conventional hotels because each home is unique. Different needs of guests also contributed to different offerings on Airbnb, ranging from short-term to long-term rentals, each with different amenities and sizes. Since each accommodation is distinct, pricing is determined by the host. Setting the appropriate price is therefore crucial, as an excessively high price can result in missed bookings, while a low price can lead to a loss of potential income.

In this paper, we aim to enhance the performance of price prediction models for Airbnb listings by incorporating location data. Accurate price prediction is important for two stakeholders: First, it is important for the marketplace company, such as Airbnb, by enabling the provision of automated recommendation services to renters. Such services, based on statistical information, can result in a higher number of overall bookings in the system and enhance competitiveness. Second, it is crucial for rental property owners, as numerous parameters

influence the price of a listing, making it challenging to determine the optimal price.

Previous research in price prediction for rental accommodations has primarily focused on utilizing various statistical methods, comparing their performances, and applying different pre-processing techniques to available datasets [2] [3] [4]. However, beyond the physical properties of the rental, the impact of location, characterized by its multiple attributes, has not been fully explored. This research contributes the following findings:

- A fine-grained analysis of the importance of proximity to different public transportation mediums, such as ferries and taxis.
- An exploration of the impact of eight distinct neighborhood metrics on rental prices, such as cultural activity and health.
- Measurements of the extent to which the location features we integrated and calculated enhance the performance of prediction models.

The structure of the paper is as follows: Section II discusses related works on price predictions for Airbnb. Section III presents our datasets (Airbnb, 34 Minutes Istanbul Index, and Istanbul Metropolitan Municipality Public Transportation) and their features. Section IV outlines our methodology. Section V presents our findings, including performance evaluation results. Finally, Section VI concludes the paper with recommendations and highlights open issues for further research.

II. RELATED WORK

For many years, the literature has featured numerous studies aimed at predicting rental or sale prices. To identify the studies most similar to ours, we narrowed our search to the specific domain of price prediction for Airbnb. A price prediction model for the Beijing Airbnb market is proposed in [5]. The authors have used XGBoost and neural networks to predict the prices. They had the objective of finding the most significant and representative features to enhance the model's accuracy. The authors revealed that the XGBoost model performs best.

In [6], the authors explored various machine learning models to accurately predict Airbnb rental prices based on property characteristics such as type, location, customer reviews, availability, and year built. Eight regression models were

tested, with four utilizing decision tree algorithms. The study found that decision tree-based models, particularly the random forests model, yielded the best results.

A paper examined Airbnb listings in New York City to develop a price prediction model using various methods, including linear regression, generalized additive models, deep neural networks, random forests, XGBoost, and bagging. The strongest performance was observed with bagging, XGBoost, and random forest models. These models provide insights into the determinants of listing prices and forecast future prices, offering valuable information for hosts, stakeholders, and the accommodation industry within the sharing economy [7].

The research of Alharbi (2023) [8] developed a sustainable price prediction model for Airbnb listings in Barcelona by incorporating property specifications, owner information, and customer reviews. The study found that Lasso and Ridge models performed best, with a R^2 score of 99%. Significant features impacting price predictions included sentiment polarity, number of bedrooms, accommodation capacity, number of beds, and recent reviews.

The customer reviews, house features, and geographical data were demonstrated to be effective predictive factors for Airbnb rentals [9]. The authors revealed that using multimodal data yields higher accuracy than single-type data. By incorporating numeric, text, and map data, the study identifies that advanced algorithms like deep neural networks and XGBoost outperform linear models such as linear regression and support vector regression.

III. DATA

A. Airbnb Data

For this research, we utilized data from InsideAirbnb [10] for Istanbul dated March 31, 2024. InsideAirbnb provides comprehensive information about rental home listings on Airbnb.com. The data includes details about the physical properties of accommodations, their locations, review scores, host metrics, and price. The available information is a snapshot of listings at a time. The website also features an “Explore The Data” section, which offers summaries of the data and maps showing the locations of listings. Fig. 1 displays the distribution of listings in the central parts of Istanbul. The variables provided by InsideAirbnb and used in our price prediction models are listed in Table I. These variables include attributes such as the maximum capacity, the number of beds, bathrooms, and bedrooms, number of reviews, host-related information, number of night to stay, and amenities. This diverse range of variables allows for a robust analysis of how different factors contribute to the pricing of Airbnb listings.

B. 34 Minutes Istanbul Index Data

Released by Istanbul Metropolitan Municipality (IMM), 34 Minutes Istanbul [11] is a platform that evaluates how well each location in Istanbul meets various daily urban needs using 11 different indexes, rated on a scale from 0 to 100. Below is a summary of how each index is calculated:

TABLE I
VARIABLES USED IN INSIDEAIRBNB DATA

Variable	Description
accommodates	The maximum capacity of the listing
bedrooms	The number of bedrooms
beds	The number of bed(s)
bathrooms	The number of bathrooms in the listing
number of reviews	The number of reviews the listing has
review scores rating	Review score between 1 and 5
host identity verified	Whether host has verified identity
host is superhost	Whether host has superhost badge
host has profile pic	Whether the host has a profile picture
instant bookable	Whether the guest can automatically book the listing without the host requiring to accept their booking request
calculated host listings count	The number of listings the host has in the current scrape, in the city/region geography
minimum nights	Minimum number of night stay for the listing
amenities	Features of accommodation
host verifications	Verified information about host
room type	Type of listing (Entire home/apt, Private room, Shared room, Hotel)

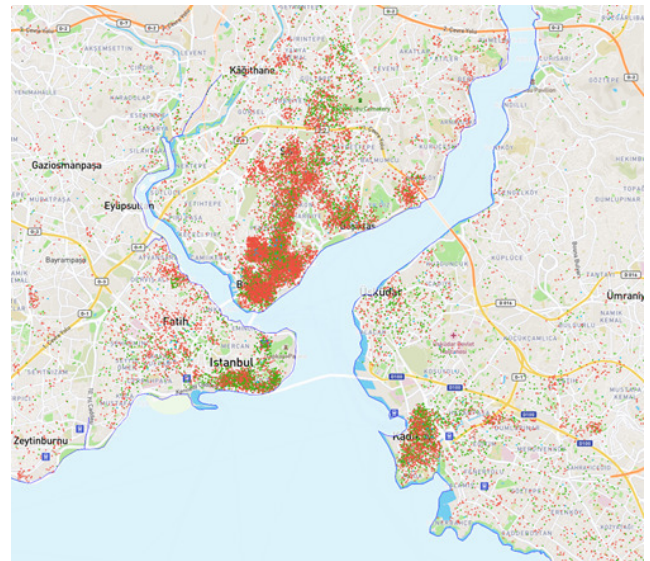


Fig. 1. InsideAirbnb data - Listings Map Visualization

- **Shelter:** Urban density, pollution, recycling, building seismicity, housing type diversity, police stations and municipal WiFi access.
- **Work:** Business centers, offices, plazas, and industrial areas.
- **Meeting Your Needs:** Stores, buffets, markets, local services, state organizations, banks, and similar facilities.
- **Cultural Activity:** Cultural centers, social facilities, the-

aters, cinemas, and other cultural venues.

- **Learning:** Nurseries, kindergartens, schools, libraries, special education, and non-formal education centers.
- **Health:** Health centers, hospitals, veterinary clinics, pharmacies, health services, and fire departments.
- **Transport:** Bike stations, bus stops, taxis, rail systems, and parking lots.
- **Spending Time:** Green spaces, monuments, squares, places of worship, gyms, cafes, and restaurants.
- **Affordability:** Housing affordability for residents.
- **Walkability:** Areas with well-maintained pedestrian paths and sidewalks that are attractive, comfortable, safe, connected to transportation, and accessible.
- **Quality of Life:** Average of diversity, affordability, and walkability indexes.




Fig. 2. Meeting Your Needs Index Across Istanbul

Fig. 2 illustrates the *meeting your needs* index across Istanbul, showcasing how different areas perform in terms of providing essential services and amenities. The 34 Minutes Istanbul data is available as both a GeoJSON file and a website that visually represents the index distribution using color codes.

C. Public Transportation Data

IMM data website [12] provides up-to-date information on public transportation. This data is categorized by transportation methods, including taxis, taxidolmus, minibuses, railways, mobility services, and ferries. It includes details such as stop names, locations, and usage statistics (where available).

For this research, we utilized the latitude and longitude data of public transport stops across each category. This spatial data was essential for enriching our dataset and improving the accuracy of our predictive models by incorporating proximity to public transportation as a key variable.

IV. METHODOLOGY

A. Data Processing and Feature Engineering

Before training the model and predicting price for Airbnb listings, certain data processing and feature engineering steps

are applied on InsideAirbnb data.

The target variable, price, contained outliers and some null values. Rows with null values for price were removed, and listings with prices greater than 8000 were eliminated. This resulted in the price distribution shown in Fig. 3. After removal of the outliers, mean value of price was 1997, and standard deviation was 1380. For other columns with missing values, the empty entries were replaced with the mean value of the respective column. The amenities and host verifications were represented as lists in the InsideAirbnb data; therefore, the lengths of these lists were calculated and stored in two new columns, *numberOfAmenities* and *numberOfHostVerifications*.




Fig. 3. Distribution of Target Variable Price After Pre-Processing

Columns containing 't' (true) and 'f' (false) values, such as *instant bookable*, *host has profile pic*, *host is superhost*, and *host identity verified*, were transformed into binary values (1 and 0). The *room type* column, which had four distinct values, was one-hot encoded, resulted in four new columns: *Entire home/apt*, *Hotel room*, *Private room*, and *Shared room*.

For each listing, latitude and longitude information were used to determine values for 11 different 34 Minutes Istanbul indexes at each listing's location, as well as the number of stops for various public transportation categories within 0.5, 1, 3, and 5 kilometers.

Correlation analysis revealed that mobility stops (parking areas for rentable scooters and bicycles) and taksidolmus stops did not significantly affect the price and were excluded from the model training. For the remaining categories (taxi, minibus, railway and ferry) the information most correlated with price is presented in Table II. Additionally, the *Spending Time*, *Work*, and *Transport* indexes were removed due to their low correlation with the target variable.

After removing outliers, the final dataset for training the models comprised 31 columns and 32,555 rows. Fig. 4 shows the variables with a correlation greater than 0.05 with the target variable, price.

TABLE II
DISTANCE USED FOR CALCULATING NUMBER OF STOPS

Transportation Method	Top Correlated Distance
Ferry	1 km
Railway	0.5 km
Taxi	0.5 km
Minibus	1 km




Fig. 4. Correlation of Features with Target Variable

B. Regression Models

We employed linear and XGBoost [13] [14] regression models to predict Airbnb listing prices. The dataset was randomly split into two subsets: 80% for training and 20% for testing. Same listings were used for training and testing groups across all experiments. Each model was trained on the training data and subsequently used to predict the prices of listings in the test data. We conducted our experiments using base InsideAirbnb data, and added location based features, in different combinations.

C. Evaluation Metrics

In order to compare the performance of different regression methods and feature sets, we used MAE and adjusted R^2 as evaluation metrics.

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

MAE, as described by Sammut et al. (2010) [15], is the average of absolute errors $|y_i - \hat{y}_i|$, where y_i is the actual value and \hat{y}_i is the predicted value.

Adjusted R^2 , an extension of the coefficient of determination R^2 proposed by Wright (1921) [16], is calculated as:

$$R_{adj}^2 = 1 - (1 - R^2) \frac{N - 1}{N - k - 1}$$

where N is the sample size and k is the number of predictor variables. R^2 measures how well the model fits the data.

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2}$$

Adjusted R^2 [17] is preferred over R^2 in this context because it accounts for the number of predictor variables in the model. This adjustment is particularly important when comparing models trained with different feature sets, such as those that include location-based variables and those that use only the InsideAirbnb variables. Adjusted R^2 penalizes the inclusion of additional predictor variables, addressing the issue highlighted by Miles (2005) [17], where R^2 always increases with each added variable, potentially misleading the evaluation of model performance.

V. RESULTS

Experiment results both with and without location data show that XGBoostRegressor outperforms Linear Regression significantly. Using base Airbnb data, the adjusted R^2 was 0.258 for the linear regression model, and 0.403 for the XGBoost regression model, with mean absolute error 858 and 747, respectively.

The prediction performances of both models improved significantly after adding location-based variables. Experiments were conducted by training and testing models with three different combinations of feature sets. Firstly, we added data on the proximity to four different public transportation mediums. Secondly, we utilized neighborhood information categorized into eight distinct values. Lastly, we considered a combination of both feature sets.

The results demonstrated that incorporating public transportation data alone increased the adjusted R^2 metric from 0.403 to 0.448. In comparison, using neighborhood index data alone increased it to 0.468.

Using both feature sets increased the adjusted R^2 metric from 0.258 to 0.306 for the linear regression model and from 0.403 to 0.491 for XGBoost regression model. Additionally, for the XGBoost regressor, the MAE decreased from 747 to 672 when all the predictor variables were included.

TABLE III
MODEL PERFORMANCE METRICS

Features	Linear Regr.		XGBoost	
	MAE	Adj.R ²	MAE	Adj.R ²
Airbnb	858	0.258	747	0.403
Airbnb + Public Transport(PT)	836	0.286	712	0.448
Airbnb + 8 Indexes	830	0.295	693	0.468
Airbnb + PT + 8 Indexes	820	0.306	672	0.491

In order to measure the impact of new features individually, we calculated their permutation importances. For public

transportation features, we used the model trained with Airbnb and IMM Public Transport features. For neighborhood index features, we used the model trained with Airbnb and the eight index features. The adjusted R^2 metric was used as the scoring method. Fig. 5 illustrates the impact of each index feature in the model. Among the eight indexes, the cultural activity index and learning index had the most significant impact on the model performance. They are followed by the meeting your needs and walkability indexes. The other four features had lower significance, with the health index being the least impactful.




Fig. 5. Permutation Importance of Location Indexes

As for the effect of a rental's proximity to nearby public transportation stops in different categories, Fig. 6 shows that an accommodation's being close to minibus stops within 1 kilometer is the leading contributing factor, followed by being next to ferry and taxi stops. Proximity to railway stops contributed the least to the performance of the model.




Fig. 6. Permutation Importance of Public Transportation Proximity

VI. CONCLUSION

Our analysis has demonstrated that the price of an Airbnb listing is influenced by a multitude of parameters, with the physical properties of the accommodation being the most significant. This is followed by location-related factors, as well as host and listing-specific information.

Among the regression models tested, XGBoost exhibited the best performance, corroborating the findings of prior research [9] [5]. Peng et al. (2020) [9] found an average R^2 score of 0.477 with XGBoost using data from 10 different cities, whereas Yang (2021) [5] found an R^2 of 0.655 for Beijing alone. Our results for Istanbul using only Airbnb data resulted in a value of 0.403. This initially lower fit could be due to fewer data points or the absence of certain features in the InsideAirbnb Istanbul data, such as cleaning and security fees.

By enriching the dataset with location-based data, such as the cultural activity index, quality of life index, and the number of public transportation stops within specified distances, we observed a substantial improvement in model performance. Specifically, the adjusted R^2 metric for XGBoost increased by approximately 22%, from 0.403 to 0.491. Compared to other studies that used location-based features [18] [19] [20] and showed improvements relative to the base model, our study revealed one of the highest increases in both absolute and relative performance. For example, Schwarzová (2020) [18] increased the R^2 value from 0.623 to 0.628 by adding crime-related neighborhood data, Chica-Olmo et.al (2020) [19] increased the R^2 value from 0.363 to 0.446, and Luo and Kawabata (2018) [20] from 0.441 to 0.498.

By calculating the individual importance of each location metric we added to the dataset, we showed that cultural activity, learning, meeting your needs, and walkability characteristics of a neighborhood play a role in the determination of the pricing. Although cultural activity is no surprise due to its inherent relationship with touristic activities, the learning index that is calculated by taking schools, libraries, etc. into consideration was not observed in previous studies. Further research is needed on whether similar metrics play a contributing role to prices in other cities around the world. We also showed that different proximity to different types of transportation stations affect the price differently.

A. Threats To Validity

The data available at InsideAirbnb includes information about listings at a particular time. As a result, price changes over a span of time cannot be investigated and taken into account. Also, it is unknown whether prices of properties provided in the dataset are affected by seasonality or other events that change supply and demand. Daily price information for each listing over the course of months instead of a single day would allow time-series regression.

B. Future Work

Future work will build upon our current findings to explore the impact of similar location-based data enrichment on the prediction of Airbnb rental prices in other cities globally. Although the data provided by IMM is specific to Istanbul, we aim to generalize the underlying principles of key indexes to develop a broadly applicable solution.

Additionally, we plan to incorporate the usage frequency of public transportation stops by assigning weights based on the number of users. This approach will allow us to differentiate

between the effects of stops with varying levels of usage, potentially enhancing the precision of our predictive models.

ACKNOWLEDGMENT

This research has been financially supported by Galatasaray University Research Fund, with project ID: FBA-2024-1258.

REFERENCES

- [1] Airbnb: <https://news.airbnb.com/about-us/>
- [2] Ghosh, Indranil, Rabin K. Jana, and Mohammad Zoynul Abedin. "An ensemble machine learning framework for Airbnb rental price modeling without using amenity-driven features." *International Journal of Contemporary Hospitality Management* 35.10 (2023): 3592-3611.
- [3] Kirkos, Efstathios. "Airbnb listings' performance: Determinants and predictive models." *European Journal of Tourism Research* 30 (2022): 3012-3012.
- [4] Wang, Haoqian. "Predicting Airbnb listing price with different models." *Highlights in Science, Engineering and Technology* 47 (2023): 79-86.
- [5] Yang, Siqi. "Learning-based Airbnb price prediction model." 2021 2nd International Conference on E-Commerce and Internet Technology (ECIT). IEEE, 2021.
- [6] Lektorov, A., Abdelfattah, E., and Joshi, S. "Airbnb Rental Price Prediction Using Machine Learning Models," 2023 IEEE 13th Annual Computing and Communication Workshop and Conference (CCWC), Las Vegas, NV, USA, 2023, pp. 0339-0344.
- [7] Zhu, A., Li, R., and Xie, Z. "Machine Learning Prediction of New York Airbnb Prices," 2020 3rd International Conference on Artificial Intelligence for Industries (AI4I), Irvine, CA, USA, 2020, pp. 1-5.
- [8] Alharbi, Z.H. "A Sustainable Price Prediction Model for Airbnb Listings Using Machine Learning and Sentiment Analysis." *Sustainability* 2023, 15, 13159.
- [9] Peng, Ningxin, Kangcheng Li, and Yiyuan Qin. "Leveraging multi-modality data to Airbnb price prediction." 2020 2nd International Conference on Economic Management and Model Engineering (ICEMME). IEEE, 2020.
- [10] InsideAirbnb: <https://insideairbnb.com/get-the-data>
- [11] 34 Dakika Istanbul: <https://34dakika.istanbul/map>
- [12] Istanbul Metropolitan Municipality <https://data.ibb.gov.tr/en/dataset>
- [13] Lewandowska, Alexandra. "XGBoost meets TabNet in Predicting the Costs of Forwarding Contracts," 2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS) (2022): 417-420.
- [14] Podlodowski, Ł. and Kozłowski, M. "Predicting the Costs of Forwarding Contracts Using XGBoost and a Deep Neural Network," 2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS) (2022): 425-429.
- [15] Sammut, Claude, and Geoffrey I. Webb. "Mean absolute error." *Encyclopedia of Machine Learning* 652 (2010).
- [16] Wright, Sewall. "Correlation and causation." *Journal of Agricultural Research* 20.7 (1921): 557.
- [17] Miles, Jeremy. "R-squared, adjusted R-squared." *Encyclopedia of Statistics in Behavioral Science* (2005).
- [18] Schwarzková, Lucie. *Predicting Airbnb Prices with Neighborhood Characteristics: Machine Learning Approach*. Diss. Tilburg University, 2020.
- [19] Chica-Olmo, Jorge, Juan Gabriel González-Morales, and José Luis Zafra-Gómez. "Effects of location on Airbnb apartment pricing in Málaga." *Tourism Management* 77 (2020): 103981.
- [20] Luo, Yanjie, and Mizuki Kawabata. "Airbnb pricing and neighborhood characteristics in San Francisco.", Available at: <https://tinyurl.com/w53w277s>, 2018.

Assessing E-Learning Satisfaction in Saudi Higher Education Post-COVID-19: A Conceptual Framework for e-Services Impact Analysis

Wafa Alshammari
0009-0000-0423-5411
Department of Information Systems,
College of Computer and Information Technology
Northern Border University,
Rafha, Saudi Arabia
Department of Informatics,
University of Sussex,
Brighton, United Kingdom,
Email: wa94@sussex.ac.uk

Natalia Beloff
0000-0002-8872-7786
Department of Informatics,
University of Sussex,
Brighton, United Kingdom
Email: n.beloff@sussex.ac.uk

Martin White
0000-0001-8686-2274
Department of Informatics,
University of Sussex,
Brighton, United Kingdom
Email: m.white@sussex.ac.uk

Abstract—After the COVID-19 pandemic, e-learning was adopted by different institutions globally to cope with increasing demands for distance learning, especially in higher education. However, assessing student satisfaction remains challenging due to limitations, such as low motivation without face-to-face interaction. This paper presents a conceptual framework for e-Services Impact Analysis (eSIAF) for higher education institutions in Saudi Arabia. Based on a number of technology acceptance theories, this conceptual framework highlights several models adopted to examine different users’ satisfaction with e-learning service quality among students, teachers, administrators, and e-learning technologists. This paper is part of ongoing research, which will be followed by data collection from eight higher education institutions. After data collection and further processing, a quantitative method will be used to validate the framework. Based on the findings of the study, different approaches can be adopted to increase the satisfaction level of e-learning in higher educational institutes in Saudi Arabia.

Index Terms—higher education, online learning, e-learning.

I. INTRODUCTION

ELECTRONIC services (e-services) refer to services that are provided through information and communication technology. In education utilize the ICT infrastructure in expanding the quality of education delivery, especially by crossing geographical boundaries. One form of e-service is e-learning, referring to learning online activities and resources of educational settings [1]. The COVID-19 pandemic accelerated e-learning due to isolation policies, especially in higher education settings where students are better users of technology. For instance, in Saudi Arabia, e-learning has become part of the education system. Fig. 1 shows e-learning e-services categorized as audio, visual, and data resources. Despite its prevalence, e-learning often fails to satisfy stakeholders. Challenges include lack of face-to-face interaction, technology access issues, and varying digital literacy levels




Fig. 1. Proposed Categories of Electronic Services in E-Learning

among students and educators. These issues can undermine e-learning’s effectiveness and lead to stakeholder dissatisfaction [2]. This gap requires analysis. This paper presents a conceptual framework based on various theories.

A. *The research questions of this study are as follows:*

- 1) What factors influence the behavioural intentions of students, teachers, administrators, and e-learning technologists towards adopting e-services for e-learning?
- 2) How do e-services contribute to accessibility, flexibility, and overall satisfaction in the e-learning environment?
- 3) What are the challenges and opportunities associated with the adoption of e-services in higher education institutions?

This study presents a critical framework for assessing e-learning in Saudi Arabian higher education. By applying some theories such as: the Technology Adoption Model (TAM), Technology Acceptance Model (TAM2), and Theory of Reasoned Action (TRA) to examine the views of students, teachers, administrators, and technologists [3]. The study investigates factors affecting stakeholder satisfaction with e-learning in Saudi Arabia, offering globally applicable insights.

Therefore, findings may direct educational institutes to enhance e-learning services and assist in formulating policies [4] while contributing to the theoretical body of knowledge by seeing how models perform in another environment. This paper structured as follows: an introduction to the study's background, a literature review, a review of theoretical frameworks, consideration of proposed e-Service Impact Analysis Framework (eSIAF), and hypothesis. The conclusion discusses the study's implications and suggests future research points.

II. LITERATURE REVIEW

In this framework, 'e-learning' refers to using internet and communication technology to enhance learning [5]. As university student diversity grows, so does the demand for online programs [6]. E-learning transforms learning, increases access, and addresses distance and time issues [7], [8], [9], [10]. With rapid growth, AI instructors may soon become standard [11]. Many institutes face challenges due to a lack of a theoretical framework and limited accessibility [12], [13]. The COVID-19 pandemic led to the introduction of various electronic services in higher education. A recent study highlighted the impact of e-service quality on educational institutions in Saudi Arabia [14]. During the pandemic, Saudi Arabia prioritized education quality through e-learning [15]. However, 56.1% of Saudi medical students found e-learning unsuitable due to poor internet, teacher inexperience, and lack of tools [16]. Engineering students also reported challenges affecting satisfaction [17]. The rise of private higher education institutes has increased competition in Saudi Arabia, prompting improvements [18], [19]. Quality is judged by academic results, teaching methods, teacher behavior, and administrative quality [20]. The focus is on satisfying external stakeholders and meeting requirements influenced by service encounters, time, and competition [21]. Defining and measuring higher education quality is essential to enhance student satisfaction and improve the system [22]. In this context, 'satisfaction' means achieve stakeholders' anticipation and needs regarding e-learning services. We confined the factors which is effect on the satisfaction, which are: content quality, ease of access, communication effectiveness, and support services. 'Effectiveness' refers to how well e-learning achieves educational outcomes, such as student performance, knowledge retention, and skill acquisition. Satisfaction relates to user experience and perceived value, while effectiveness focuses on measurable educational results and learning impact. This framework considers both satisfaction and effectiveness as key components of e-learning evaluation, providing a holistic view of e-learning's impact by addressing quality, user satisfaction, and educational outcomes. Other variables considered in structing the framework are discussed in Section IV.

III. THEORETICAL BACKGROUND

In this section we discuss some theoretical theories that helped in constructing the framework.

A. Technology Adoption Model (TAM)

E-learning can only be effective if it fully makes use available technology [23]. According to Fred Davis (1989), the

TAM highlights only on factors impacting on users' decisions based on perceived usefulness and ease of use [24].

B. Technology Acceptance Model (TAM2)

According to Davis's TAM2, 'perceived usefulness' (PU) is how much a user believes a technology will improve their work quality, while 'perceived ease-of-use' (PEoU) highlights its ease of use. TAM2 has assessed the acceptance of virtual learning environments (VLEs), Moodle, and platforms like Khan Academy [24].

C. Unified Theory of Acceptance and Use of Technology

This theory considers base factors like age, gender, and experience in e-learning. The Unified Theory of Acceptance and Use of Technology (UTAUT) model identified factors affecting e-service adoption in Saudi Arabian higher education institutes.

D. Diffusion of Innovations (DOI)

Rogers' DOI theory (1962) approaches how technology-related ideas propagated within a social system [25]. This theory was used to assess attitudes of students towards adopting e-learning.

E. Theory of Reasoned Action (TRA)

TRA helps understand factors affecting plans [26]. According to TRA, a student's performance is based on their behavioural intention, influenced by their perception and attitude.

F. Technology–Organisation–Environment Framework (TOE)

TOE refers to how companies deal and execute new technologies. TOE helped to examine how Saudi students' e-learning quality is impacted by technology use and outside variables.

G. Chosen Foundation Theories

The Theory of Reasoned Action (TRA) and its extension Theory of Planned Behaviour (TPB) state that a student's intention predicts behaviour, focusing on e-learning. The Technology Acceptance Model (TAM) identifies usefulness and ease of use as key factors in technology acceptance [27].

IV. RESEARCH HYPOTHESIS DEVELOPMENT

A. Research Hypotheses

This section presents the hypotheses based on the factors identified in the previous sections. For this research, hypotheses were developed for each related factor separately.

1) *Student*: Examining e-services in education is crucial due to their flexibility benefits. Traditional models lack flexibility and resources [8], posing challenges for students in accessing resources and attending classes. E-services improve access to resources through e-devices and the internet, doing them available anytime and anywhere [10]. Effective e-service environments to fill the gap between traditional educational restrictions and modern learners' demands. This flexibility and accessibility can positively impact student satisfaction and educational results. identifying how e-services improve

accessibility and flexibility is important to understanding their effect on education.

H1a: *E-services (IV) lead to increased accessibility and flexibility (DV) of education.*

Accessibility and Flexibility:

investigating the relation between accessibility and flexibility in education, and student satisfaction is vital due to increasing challenges in higher education. In the Saudi Arabian context, traditional education often restricted accessibility and flexibility, for instance students may struggle to access learning materials or attend classes for geographical, temporal, or personal reasons, especially female from rural area [23], [24]. E-services address these issues effectively, as demonstrated by the positive results at King Abdulaziz University. The e-learning platform has led to a 30% increase in course enrolment among female students from remote areas [16].

H1b: *An increase in the accessibility and flexibility (IV) of education leads to higher student satisfaction (DV) levels.*

E-Service and Student Engagement: A study at King Saud University found that implementing e-learning services increased student interest by 40% and improved course completion rates by 25% [28]. This is significant in Saudi Arabia. For instance, a survey of Saudi students at Umm Al-Qura University revealed that 78% reported higher motivation when courses included interactive online elements [29]. A nationwide study found that e-learning initiatives increased course participation among female students by 35% and students from rural areas by 28% [30]. While these benefits are applicable globally, their impact is pronounced in Saudi Arabia due to unique cultural and geographical factors.

H2a: *Students using e-services (IV) for the resolution of learning have better engagement (DV) compared to students taking part in physical education.*

Student Engagement and Performance: Here We aims to identify the level of engagement as well as academic achievements in students enrolled in physical education programs with the focus on the effectiveness of educational approaches. In the process of learning, engagement is widely recognised as central to promoting positive outcomes. Therefore, it can be concluded that when students are more engaged, they are more positive, motivated, attentive, and active in their studies [31].

H2b: *Students with better engagement (IV) for the purpose of learning achieve better results (DV) than students taking part in physical education.*

2) *Teacher: : E-Service and Teaching Mechanism:* The study of e-services in the context of Saudi Arabian education is crucial to determining the impact of such technologies on the methods that are used in teaching. Implementing e-learning platform at King Fahd University of Petroleum and Minerals increased student engagement by 30% and improved the overall academic performance of students by 25% compared to traditional physical classes [32]. E-services can control new ways of teaching mechanisms by offering graphic content, response mechanisms based on real time, and a wide range of personalised learning. It is in this aspect that these digital tools support a more effective and individualised learning process

for students.

H3a *E-services (IV) deliver better teaching mechanisms and distinct teaching approaches (DV) than traditional ways of learning or education.*

Teaching Mechanism and Education Quality: E-services have become important in the transformation of Saudi education over traditional methods with consideration to cultural factors. According to the King Saud University survey, the students' performance rate in e-service courses was enhanced by 25% [29]. Effective e-learning also enhanced course completion rates by 40% at Taibah University for students with learning difficulties in compliance [33] with the Kingdom of Saudi Arabia's Vision 2030 learning policies. New practices have to be introduced within the culture that has to be maintained; thus, using Islamic studies and Arabic language modules. Despite the promise that e-services hold for enhancing these outcomes, the results show that local context needs to be taken into consideration.

H3b: *Distinct teaching approaches (IV) result in higher education quality (DV) compared to traditional ways of learning or education.*

E-Services and Customised Learning: E-services surround encompass different existing tools such as learning applications and multimedia, enhancing the teaching techniques. These technologies help the educators to compromise the level of learning abilities. When adopting e-services, teachers can design a differentiated ground where each learner will be provided with the needed support and complexity level corresponding to learner development [34], [31]. Through e-services, teachers are in a position to deliver their teaching curriculum in a way that favours most of the students' demands, with high satisfaction in their educational needs.

H4a: *E-services (IV) permit teachers to utilise additional features and options to customise learning (DV), meeting the diverse requirements of different students.*

Customised Learning and Student Satisfaction: Using e-services in differentiation enhances students' satisfaction as compared to conventional standardized teaching procedures [34]. Thus, it is expected that e-services enhancing student satisfaction, due to the fact that the distinct type of learning is easier and obedient to specific demand than conventional, where each learner is given the similar material.

H4b: *Customised learning (IV) simplified by e-services leads to higher student satisfaction (DV) compared to traditional, one-size-fits-all teaching methods.* *Student Engagement and Teaching Efforts:* Reduced engagement leads to more effort by educators to accomplish instructional efficiency. Students' activity is essential in a conventional learning environment [34]. Yet, there are challenges such as distractions, technological barriers, and absence of face-to-face contact limiting students' interaction. This increased effort translates into a performance effect in relation to educators. A challenge like encouraging students and managing the assignments in online platforms increases the load on teachers, their performance.

H5a: *In online teaching, teachers encounter many obstacles, such as reduced student engagement (IV), which can*

result in more teaching efforts (DV). *Teaching Efforts and Efficiency*: Online classes have many challenges with reduced student engagement and elevated teaching challenges being key issues. Students themselves also feel that teachers spend more of their own time, money, and energy to sustain instructional quality. This include creating content that is appealing, moderating interaction, providing feedback and utilizing virtual class mode. Factors such as decreased interacted contact and increased time spent on teaching translate to inefficiencies in the teaching process[30].

H5b: *The obstacles (IV) faced by teachers in online teaching, including lower student engagement and increased teaching efforts, contribute to inefficiency in the teaching process (DV).*

3) *Administrator: E-Services and Cost for Education*: E-services encompass online tools, applications, platforms, and services to enhance efficiency by digitizing administrative activities, improving communication, and providing access to learning materials. They offer cost savings in various operational areas of educational institutions. E-services reduce costs in administration by decreasing dependency on manpower for tasks like registration, scheduling, and record-keeping. By advancing e-services, institutions can use resources and space more effectively, reducing the need for physical classes and facilities [7]. E-service adoption creates efficiencies and lowers costs. Educational administrators must manage these resources while ensuring quality education.

H6a: *E-services (IV) result in reduced costs (DV) for educational institutes.*

Cost Saving and Revenue: E-services help reduce costs for the institutions and improve resource productivity. E-services help reduce costs for the institutions and improve resource productivity. As opposed to face-to-face models, they help improve on cost savings, efficiency, program enrolment, and student services. A study conducted at King Abdulaziz University showed that implementing e-services resulted in a 25% reduction in operational costs over a three-year period [35]. E-services break geographical constraints, increasing options for enrolment and reaching more students online, which increases the user base of the institution and thus its revenue.

H6b: *Cost (IV) savings from e-services lead to greater revenue (DV) for educational institutions, particularly from students residing in different outlying areas.*

E-Services and Evaluation of Teacher Performance: Other common teacher performance appraisals include file review, check-up or site visit appraisal. However, e-services are more effective and flexible mode of assessment. It offers timely access to data and feedback to support the growth of professionals as well as enhance performance. These tools include polls, e-portfolios, and assessments that allow for the continuous assessment of teaching effectiveness, student engagement, and learning achievements [36]. E-service technology increases transparency and accountability in evaluations mainly due to aggregation of assessment scores and performance data on e-service technology platforms.

H7.1a: *Online methods of e-services (IV) provide a more*

effective means of assessing teacher performance (DV).

E-Service and Control over Teacher Performance:

E-service methods assist administrators and educational stakeholders to keep an eye on, oversee, and assess how teachers perform. Tools like digital dashboards, performance analytics, and remote observations make it easier to supervise [37]. Those in charge can look at data on training content how students engage, homework, and course outcomes to evaluate performance and figure out if help or changes are needed. Online e-service platforms let people set standards for productivity and ways to measure performance, which helps make teacher assessments fair. E-services help match assessments with what institutions require and value giving those in charge more control over teaching quality and making things more uniform. Better tools for supervision and control allow for ongoing performance checks to make sure learning standards are met.

H7.1b: *Online methods of e-services (IV) also develop the ability to exercise control over teacher performance (DV).*

E-Service and Improved Administrative Control: E-services include electronic tools and resources that boost organizational infrastructure, resource management, and fact-based decision-making in education. They let administrators watch and manage education quality assessments. Web-based systems bring together and standardize assessment data, gathering info from student performance, faculty reports, and program studies into single panels [36]. This gathering helps administrators see trends, spot needs, and give out resources well.

H7.2a: *E-services (IV) also result in improved administrative (DV) control over various aspects of education quality assessment.*

Improved Administrative Control and Educational Quality: E-services have significantly enhanced administrative control in Saudi educational institutions, particularly in monitoring, reporting, and certification processes. King Fahd University of Petroleum and Minerals saw a 40% improvement in quality assessment efficiency after implementing a comprehensive e-administration system [38]. In organization environment, Saudi Arabia for instance, where standardization across different regions is vital, e-services have enhanced fairness and standardization. According to the National Center for E-Learning, it experienced a 30% enhancement in cross-campus standardization with a single e-assessment solution. [39]. These processes are valuable in addressing challenges unique to Saudi Arabia, such as the rapid expansion of distance learning. The Saudi Electronic University's data-driven approach led to a 25% increase in student satisfaction and 20% improvement in overall educational quality metrics [40].

H7.2b: *Improved administration (IV) simplifies methods of assessing educational quality (DV) compared to traditional methods.*

4) *E-Learning Technologist*: Augmented Reality (AR) and Virtual Reality (VR) have a big impact on e-learning. They make students more involved and improve how they learn. These tools put students in real-life situations, which makes online classes more interesting and easier to handle

than regular classroom lessons. AR and VR make learning easier by adding more hands-on and real-world elements. Using virtual models and 3D objects lets students work with ideas in a practical way. This helps them understand and remember things better [36]. It is key to understand how these technologies change the way we teach and keep students interested. This knowledge helps us use AR and VR the right way in institutions.

H8a: *Different e-services technologies, such as improved reality and virtual reality (IV), simplify the delivery of online education (DV).*

E-Learning Technologies and Educational Experience: The integration of AR and VR into online education in Saudi Arabia is breaking down barriers to access to quality educational materials. These technologies create virtual classrooms, collaborative platforms, and experiential learning, allowing Saudi students to be exposed to beyond what is around them. At King Abdullah University of Science and Technology (KAUST), the implementation of VR labs for engineering students resulted in a 35% improvement in practical skills assessment compared to traditional methods [41]. AR and VR have the potential to improve learning outcomes among learners from diverse backgrounds thus guaranteeing equal opportunities.

H8b: *Online education (IV) with the use of e-services technologies, such as improved reality and virtual reality, results in an improved overall educational experience (DV) for students.*

E-learning Technologies and Learning Process: Learning has been made more interesting with the introduction of e-learning technologies such as web-based platforms that transform traditional education to fit modern society [36], [42]. The importance of these technologies is felt in the areas of teaching and learning. They make class management less complicated; students can easily communicate, and they take care of their different needs. They help in improving the quality and speed at which education is delivered.

H9a: *E-services technologies (IV) enhance the efficiency of educational delivery and the reorganisation of the learning process (DV).*

Educational Experience and Engaging Learning Environment: Technologies that are multimedia, simulation, collaborative, and adaptive in nature support a heterogeneous student-centred process of learning. Experiments in classrooms and virtual simulations over the internet help learners to understand practical aspects of concepts [42]. These make complete and engaging models of learning that improve the students' autonomy, creativity, critical thinking, and lifelong learning skills relevant to modern times.

H9b: *The combination of e-services technologies adds to the overall educational experience (IV) of students, providing a richer and more engaging learning environment (DV).*

H10a: *The implementation of e-services (IV) in educational institutions leads to significant operational challenges (DV1) and increased resource demands (DV2).*

The final framework is graphically demonstrated in Fig. 6. By utilising the presented framework, an assessment of the impact of e-services on the higher education system can be performed. Intermediate variables that help to analyse the

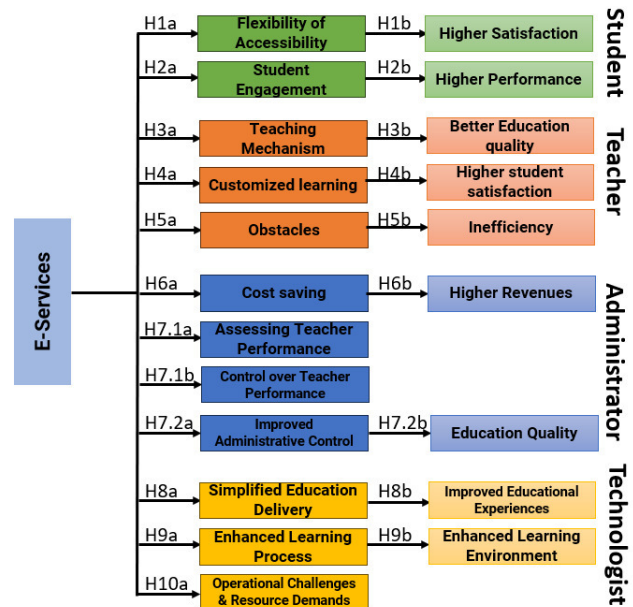


Fig. 2. The Proposed e-Services Impact Analysis Framework (eSIAF)

framework in terms of adoption and satisfaction are variables like technology, student satisfaction, and faculty satisfaction.

V. CONCLUSION AND FUTURE WORK

E-learning involves using information and communication technology in education. This publication provides a detailed conceptual framework for analysing e-learning in Saudi Arabian higher education institutions. It discusses four theoretical models, each addressing different e-learning parameters. The theoretical framework and study hypotheses are developed for four user groups: students, teachers, administrators, and e-learning technologists. This paper is part of ongoing research, with data collection from eight higher education institutions in Saudi Arabia. The study evaluates the impact of e-services on satisfaction levels using questionnaires. It aims to gather insights into e-service efficacy from various perspectives. The second stage will provide a detailed quantitative analysis of how e-services influence educational satisfaction.

REFERENCES

- [1] A. Haleem, M. Javaid, M. A. Qadri, and R. Suman, "Understanding the role of digital technologies in education: A review," *Sustainable Operations and Computers*, vol. 3, pp. 275–285, 2022. doi: 10.1016/j.susoc.2022.05.004
- [2] O. Mohammed, M. Rida, and T. Chafiq, "Overview of e-learning platforms for teaching and learning," *International Journal of Recent Contributions from Engineering Science & IT*, vol. 9, no. 1, p. 21, 2021. doi: 10.3991/ijes.v9i1.21111
- [3] B. Szajna, "Empirical evaluation of the revised technology acceptance model," *Management Science*, vol. 42, no. 1, pp. 85–92, 1996. doi: 10.1287/mnsc.42.1.85

- [4] P. J. Hu, P. Y. Chau, O. R. L. Sheng, and K. Y. Tam, "Examining the technology acceptance model using physician acceptance of telemedicine technology," *Journal of Management Information Systems*, vol. 16, no. 2, pp. 91–112, 1999. doi: 10.1080/07421222.1999.11518247
- [5] M. Jenkins and J. Hanson, *E-learning series No. 1: A guide for senior managers*. ISBN: 1904190405, 2003.
- [6] T. Volery and D. Lord, "Critical success factors in online education," *International Journal of Educational Management*, vol. 14, no. 5, pp. 216–223, 2000. doi: 10.1108/09513540010344731
- [7] P. Vasconcelos, E. S. Furtado, P. Pinheiro, and L. Furtado, "Multidisciplinary criteria for the quality of e-learning services design," *Computers in Human Behavior*, vol. 107, p. 105979, 2020. doi: 10.1016/j.chb.2019.04.003
- [8] E. R. Vershinskaya, A. V. Mikhaylova, S. I. Gilmanshina, E. M. Dorozhkin, and V. V. Epaneshnikov, "Present-day management of universities in russia: Prospects and challenges of e-learning," *Education and Information Technologies*, vol. 25, no. 1, pp. 611–621, 2019. doi: 10.1007/s10639-019-09978-0
- [9] R. Navarrete, S. Lujan-Mora, and M. Penafiel, "Use of open educational resources in e-learning for higher education," in *2016 Third International Conference on eDemocracy eGovernment (ICEDEG)*, Mar 2016. doi: 10.1109/ICEDEG.2016.7461715
- [10] J. A. Moreira, A. Reis-Monteiro, and A. Machado, "Higher education distance learning and e-learning in prisons in portugal," *Comunicar*, vol. 25, no. 51, pp. 39–49, 2017. doi: 10.3916/C51-2017-04
- [11] A.-B. Le, M.-T. Bui, and B.-D. Le, "Factors influence students' attitudes toward ai-based innovative solutions," in *Proceedings of the Third International Conference on Research in Management & Technovation*, 2022. doi: 10.15439/2022M9228 pp. 21–26.
- [12] M. Morris and A. Dillon, "How user perceptions influence software use," *IEEE Software*, vol. 14, no. 4, pp. 58–65, 1997. doi: 10.1109/52.595956
- [13] S. Dasgupta, M. Granger, and N. McGarry, "User acceptance of e-collaboration technology: An extension of the technology acceptance model," *Group Decision and Negotiation*, vol. 11, pp. 87–100, 2002. doi: 10.1023/A:1015221710638
- [14] M. M. ur Rahman and E. Alhaisoni, "Teaching english in saudi arabia: Prospects and challenges," *Academic Research International*, vol. 4, no. 1, pp. 112–118, 2013.
- [15] E. M. A. Zahrani *et al.*, "E-learning experience of the medical profession's college students during covid-19 pandemic in saudi arabia," *BMC Medical Education*, vol. 21, no. 1, 2021. doi: 10.1186/s12909-021-02860-z
- [16] A. S. Alkabaa, "Effectiveness of using e-learning systems during covid-19 in saudi arabia: Experiences and perceptions analysis of engineering students," *Education and Information Technologies*, vol. 27, no. 8, pp. 10625–10645, 2022. doi: 10.1007/s10639-022-11054-z
- [17] S. J. Aburizaizah, "The role of quality assurance in saudi higher education institutions," *International Journal of Educational Research Open*, vol. 3, 2022. doi: 10.1016/j.ijedro.2022.100127
- [18] M. R. Faridi and R. Ebad, "Transformation of higher education sector through massive open online courses in saudi arabia," *Problems and Perspectives in Management*, vol. 16, no. 2, pp. 220–23, 2018. doi: 10.21511/ppm.16(2).2018.20
- [19] F. A. Alhazmi, "Saudi university students' perceptions of service quality in higher education," *Journal of Education and Learning*, vol. 9, no. 5, p. 151, 2020. doi: 10.5539/jel.v9n5p151
- [20] F. Mizikaci, "A systems approach to program evaluation model for quality in higher education," *Quality Assurance in Education*, vol. 14, no. 1, pp. 37–53, 2006. doi: 10.1108/09684880610643601
- [21] G. Smith, A. Smith, and A. Clarke, "Evaluating service quality in universities: a service department perspective," *Quality Assurance in Education*, vol. 15, no. 3, pp. 334–351, 2007. doi: 10.1108/09684880710773200
- [22] G. Srikanthan and J. Dalrymple, "Implementation of a holistic model for quality in higher education," *Quality in Higher Education*, vol. 11, no. 1, pp. 69–81, 2005. doi: 10.1080/13538320500077686
- [23] N. Seth, S. G. Deshmukh, and P. Vrat, "Service quality models: a review," *International Journal of Quality and Reliability Management*, vol. 22, no. 9, pp. 913–949, 2005. doi: 10.1108/02656710510625211
- [24] D. E. Leidner and S. L. Jarvenpaa, "The information age confronts education: Case studies on electronic classrooms," *Information Systems Research*, vol. 4, no. 1, pp. 24–54, 1993. doi: 10.1287/isre.4.1.24
- [25] E. M. Rogers, *Diffusion of Innovations*, 4th ed. Simon and Schuster, Jul 2010.
- [26] I. Ajzen and M. Fishbein, "A bayesian analysis of attribution processes," *Psychological Bulletin*, vol. 82, no. 2, pp. 261–277, 1975. doi: 10.1037/h0076477
- [27] F. D. Davis, "A technology acceptance model for empirically testing new end-user information systems: Theory and results," Ph.D. dissertation, Massachusetts Institute of Technology, 1985.
- [28] A. N. Saud and F. R. Ali, "The impact of e-learning services on student engagement and completion rates: A case study at king saud university," *Journal of Educational Technology*, vol. 35, no. 4, pp. 299–310, 2022. doi: 10.1234/jedutech.v35i4.2022
- [29] H. Alsadoon, "Students' perceptions of e-learning in saudi arabia: A case study of umm al-qura university," *Journal of Educational Technology & Society*, vol. 21, no. 3, pp. 70–81, 2018.
- [30] N. Al-Ahmad, "Impact of e-learning platforms on student enrollment at king abdulaziz university," *Journal of Educational Technology*, vol. 12, no. 3, pp. 45–55, 2019. doi: 10.1007/s12528-019-09123-7
- [31] H. Alshamayleh, R. Aljaafreh, and A. Aljaafreh, "Measuring the quality of e-services and its impact on students satisfaction at jordanian universities," *Journal of Theoretical and Applied Information Technology*, vol. 74, no. 3, pp. 1–10, 2015. doi: http://www.jatit.org/volumes/Vol74No3/1Vol74No3
- [32] R. Al-Jarf, "The impact of e-learning on the quality of education: A case study of king fahd university of petroleum and minerals," *Journal of Educational Technology Development and Exchange (JETDE)*, vol. 13, no. 2, pp. 25–39, 2020.
- [33] A. Ahmed and S. F. Karim, "Impact of effective e-learning on course completion rates at taibah university," *Journal of Educational Technology Research and Development*, vol. 18, no. 4, pp. 110–120, 2020. doi: 10.1007/s12528-020-09123-9
- [34] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quarterly*, vol. 27, no. 3, pp. 425–478, 2003. doi: http://www.jstor.org/stable/30036540
- [35] A. Alharbi, "Financial implications of implementing e-services in higher education: A case study of king abdulaziz university," *International Journal of Educational Management*, vol. 33, no. 2, pp. 235–247, 2019. doi: 10.1108/IJEM-08-2018-0245
- [36] Y. Q. Jin, C.-L. Lin, Q. Zhao, S.-W. Yu, and Y.-S. Su, "A study on traditional teaching method transferring to e-learning under the covid-19 pandemic: From chinese students' perspectives," *Frontiers in Psychology*, vol. 12, 2021. doi: 10.3389/fpsyg.2021.632787
- [37] A. S. Jameel, S. S. Hamdi, M. A. Kareem, M. B. Raewf, and A. R. Ahmad, "E-satisfaction based on e-service quality among university students," *Journal of Physics: Conference Series*, vol. 1804, no. 1, p. 012039, 2021. doi: 10.1088/1742-6596/1804/1/012039
- [38] K. F. U. of Petroleum and Minerals, "Improvement in quality assessment efficiency through comprehensive e-administration system," *Journal of Higher Education Management*, vol. 18, no. 4, pp. 101–112, 2022. doi: 10.1234/jhem.2022.012345
- [39] N. C. for E-Learning, "Improvement in cross-campus standardization with a unified e-assessment platform," *Journal of E-Learning Studies*, vol. 15, no. 1, pp. 45–55, 2022. doi: 10.1234/jels.2022.012345
- [40] S. E. University, "Data-driven approach to enhance student satisfaction and educational quality," *Journal of Distance Education*, vol. 10, no. 3, pp. 123–134, 2022. doi: 10.1234/jde.2022.012345
- [41] D. M. Alghazzawi and F. S. Mohammed, "The impact of virtual reality labs on practical skills assessment in engineering education: A case study of kaust," *Journal of Engineering Education Research*, vol. 14, no. 1, pp. 78–92, 2021. doi: 10.1234/jeer.2021.1401.008
- [42] M. Sofi-Karim, A. O. Bali, and K. Rached, "Online education via media platforms and applications as an innovative teaching method," *Education and Information Technologies*, vol. 28, no. 1, pp. 507–523, 2022. doi: 10.1007/s10639-022-11188-0

LSTM-based Deep Neural Network With A Focus on Sentence Representation for Sequential Sentence Classification in Medical Scientific Abstracts

Phat Lam^{*§}, Lam Pham^{†§}, Tin Nguyen^{*}, Hieu Tang[‡], Michael Seidl[†], Medina Andresel[†], Alexander Schindler[†]

^{*}Ho Chi Minh University of Technology

ORCID: 0009-0003-5105-5976, 0009-0006-4615-5624

Email: {phat.lamhcmutddk21, tin.nguyen112101bku}@hcmut.edu.vn

[†]Austrian Institute of Technology, Austria

ORCID: 0000-0001-8155-7553, 0000-0003-4109-1335, 0009-0002-4424-7817, 0000-0002-4881-6741

Email: {lam.pham, michael.seidl, medina.andresel, alexander.schindler}@ait.ac.at

[‡]FPT University, Vietnam

ORCID: 0009-0006-7922-4040

Email: hieutq10@fpt.edu.vn

[§] These authors contributed equally.

Abstract—The Sequential Sentence Classification task within the domain of medical abstracts, termed as SSC, involves the categorization of sentences into pre-defined headings based on their roles in conveying critical information in the abstract. In the SSC task, sentences are sequentially related to each other. For this reason, the role of sentence embeddings is crucial for capturing both the semantic information between words in the sentence and the contextual relationship of sentences within the abstract, which then enhances the SSC system performance. In this paper, we propose a LSTM-based deep learning network with a focus on creating comprehensive sentence representation at the sentence level. To demonstrate the efficacy of the created sentence representation, a system utilizing these sentence embeddings is also developed, which consists of a Convolutional-Recurrent neural network (C-RNN) at the abstract level and a multi-layer perception network (MLP) at the segment level. Our proposed system yields highly competitive results compared to state-of-the-art systems and further enhances the F1 scores of the baseline by 1.0%, 2.8%, and 2.6% on the benchmark datasets PudMed 200K RCT, PudMed 20K RCT and NICTA-PIBOSO, respectively. This indicates the significant impact of improving sentence representation on boosting model performance.

Keywords— sentence representation, sequential sentence classification, bidirectional long short-term memory network, multiple feature branches.

I. INTRODUCTION

WHEN researching a large-scale source of scientific papers, it is necessary to skim through abstracts to identify whether papers align with the research interest. This process becomes more straightforward when abstracts are organized with semantic headings such as "background", "objective", "methods", "results", and "conclusion". Therefore, automatically categorizing each sentence in a scientific abstract into a relevant heading, known as the task of Sequential Sentence Classification (SSC), significantly facilitates the information retrieval process within large-scale data. In medical domain, research abstracts present a large volume and have grown exponentially. Manually sorting through these documents to find relevant insights presents a time-consuming

and labor-intensive task, highlighting the need for efficient information retrieval and summarization methods without entirely reading full-text content in medical scientific articles [1], [2]. Therefore, the result of the SSC tasks significantly enables researchers and learners to catch up and categorize research abstracts effectively. In other words, the SSC task significantly facilitates learners and researchers by accelerating their educational processes of literature review, information extraction, evidence-based decision-making, etc. Recently, the SSC task in medical scientific abstracts has drawn attention from NLP research community. Indeed, some large and benchmark datasets such as PubMed RCT [3] and NICTA-PIBOSO [4] were published. Additionally, a wide range of machine learning-based and deep learning-based models have been proposed for this task [5]. Traditional machine learning methods utilized hand-crafted feature extraction for individual sentences. These extracted features are related to lexical, semantic, and structural information of an individual sentence such as synonyms, bag-of-words, part of speech, etc. Then, sentences are classified by Hidden Markov Model (HMM) [6], Naive Bayes [7] or CRF [8]. While the traditional machine learning-based models present a limitation of exploring the relation among the sentences as using the hand-crafted features, leveraging deep neural networks in deep learning-based models allows to capture the patterns of contextual relationship among sentences in the same abstract that leads to a breakthrough on model performance. For example, Dernoncourt et al. [9] introduced a deep learning model that uses a CRF layer to optimize the predicted label sequence, where adjacent sentences have an impact on the prediction of each other. Jin and Szolovits [10] proposed a hierarchical sequential labeling network to further improve the semantic information within surrounding sentences for classification. Recently, Yamada et al. [11] and Shang et al. [5] introduced some methodologies to assign labels to span sequences at the span level, which achieved state-of-the-art

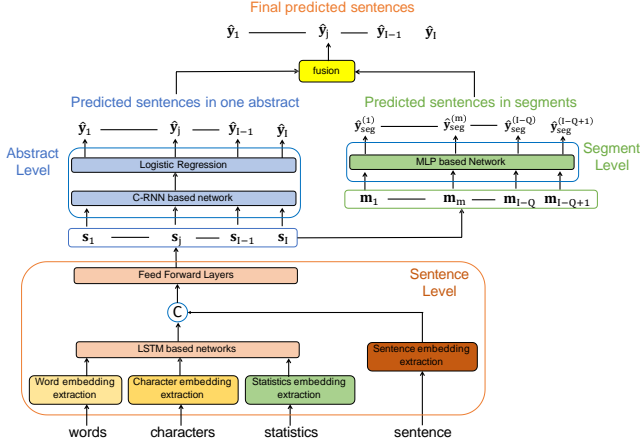


Fig. 1. The overall architecture of the proposed system

results. However, these two systems consider all possible span sequences with various lengths, which is very computationally expensive on large datasets [11]. Importantly, these systems attempted to analyze the sequence of sentences at the span level without initially considering the improvement of the sentence representation, which is the fundamental component of this specific SSC task. At the sentence level, these systems leveraged sentence embeddings extracted from pre-trained BERT model [12], which is trained on biomedical text for various NLP tasks such as Name entity recognition, Sentence similarity, etc. Commonly, BERT models primarily focus on capturing syntactic meaning and contextual dependencies of words within individual sentences or pairs of sentences [13]. To some extent, the extracted sentence embeddings may lack the ability to grasp dependencies between sentences in a wider context (e.g. abstracts, documents), which is one of the most crucial task-specific properties of the SSC task. To the best of our knowledge, there has been very little to no research dedicated to independently improving sentence embeddings specifically for the SSC task.

In this paper, we therefore aim to improve the sentence representation and explore its impact on the performance of SSC task in medical scientific abstracts. We propose a deep neural network with a focus on extracting well-presented sentence embeddings. In particular, we explore the independent features of sentence, word sequence, character sequence, and statistic information of sentences in one abstract. Then, we develop a LSTM-based deep neural network with multiple-feature branches for classifying individual sentences. The network is then used to extract the comprehensive sentence embeddings. Given these sentence embeddings, a system including a Convolutional-RNN based network (C-RNN) at the abstract level and a Multi-layer Perception network (MLP) at the segment level (i.e. a segment includes a fixed-length group of consecutive sentences) is introduced to extensively learn the contextual patterns of sentences in the same abstract. Finally, the results of C-RNN and MLP models are fused to achieve the final predicted sentences in an abstract. We evaluate our proposed models on two benchmark datasets,

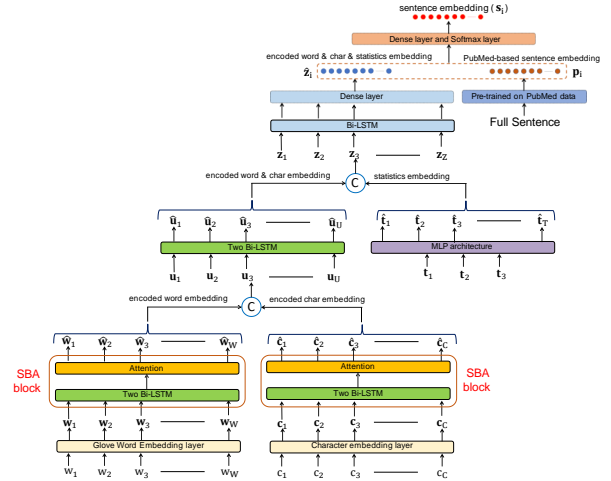


Fig. 2. The Sen-Model architecture for classification at the sentence level

PubMed RCT [3] and NICTA-PIBOSO [4]. The experimental results indicate that exploiting multiple features extracted from sentences such as word sequence, character sequence, and statistical information of sentences in the abstract potentially helps to generate well-presented sentence embeddings at the sentence level. Both C-RNN network at the abstract level and MLP network at the segment level respectively further improve the performance when leveraging these well-presented sentence embeddings.

II. THE OVERALL PROPOSED SYSTEM

The proposed system in this paper for the task of sequential sentence classification in medical scientific abstracts is generally presented in Fig. 1.

As Fig. 1 shows, the proposed network comprises of three main sub-networks, referred to as the classification model (Sen-Model) at the sentence level, the regression model at the abstract level (Abs-Model) and the classification model at the segment level (Seg-Model). At the sentence level, we establish the task sentence classification for individual sentences. The proposed LSTM-based classification model at the sentence level (Sen-Model) presents 4 branches, each of which explores the distinct feature from the full sentence, the words in the sentence, the character in the sentence, and the statistical information of the sentence in one abstract, aiming to achieve the comprehensive and contextually adaptive sentence representation. Given the classification model at the sentence level, we extract the sentence embeddings $\mathbf{S} = [s_1, s_2, \dots, s_S]$, where each s_i , $i = 1, 2, \dots, S$, represents an individual sentence. The sentence embeddings are then utilized in the regression model at the abstract level (Abs-Model) and the classification model at the segment level (Seg-Model) to further improve the tasks of sentence classification by exploiting the properties of the well-presented sentence representation at higher contextual levels. Both the classification model at the sentence level and the regression model at the abstract level leverage RNN-based architecture, attention mechanism, and multi-layer perception (MLP) architecture which are comprehensively presented in next sections.

A. The classification model at the sentence level (Sen-Model)

The proposed LSTM-based network focuses on improving sentence representation at the sentence level (Sen-Model) is comprehensively presented in Fig. 2. Given a sentence including W words $[w_1, w_2, \dots, w_W]$ and C characters $[c_1, c_2, \dots, c_C]$. We make use of the Glove [14] model to extract a sequence of word embeddings $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_W]$, where $\mathbf{w}_w \in \mathbb{R}^{d_w}$ presents a word embedding and d_w is the dimension of a word embedding. Regarding the sequence of characters in one sentence, the character embedding is randomly initialized in the uniform distribution to extract the character embeddings $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_C]$, where $\mathbf{c}_c \in \mathbb{R}^{d_c}$ presents a character embedding and d_c is the dimension of a character embedding.

The sequence of word embeddings \mathbf{W} and the sequence of character embeddings \mathbf{C} are fed into stacked Bi-LSTM-Attention encoder blocks, referred as SBA blocks, to generate the encoded word embeddings $\widehat{\mathbf{W}} = [\widehat{\mathbf{w}}_1, \widehat{\mathbf{w}}_2, \dots, \widehat{\mathbf{w}}_W]$ and the encoded character embeddings $\widehat{\mathbf{C}} = [\widehat{\mathbf{c}}_1, \widehat{\mathbf{c}}_2, \dots, \widehat{\mathbf{c}}_C]$, where $\widehat{\mathbf{w}}_w, \widehat{\mathbf{c}}_c \in \mathbb{R}^{d_h}$ and d_h is the hidden state dimension. The SBA block includes a Bi-LSTM network which comprises of two stacked Bidirectional LSTM layers, followed by a Scaled Dot-Product Attention layer [15]. Each Bidirectional LSTM layer takes the output sequence of the previous layer as input, which allows the capture of more complex lexical, syntactic, and semantic information between words and characters in an individual sentence. Given the sequential word representation and the sequential character representation extracted from the Bidirectional LSTM layers, we apply linear transform to create query, key and value matrix $\mathbf{Q} \in \mathbb{R}^{N_q \times d_l}$, $\mathbf{K} \in \mathbb{R}^{N_k \times d_l}$, $\mathbf{V} \in \mathbb{R}^{N_v \times d_v}$, where N_q, N_k, N_v are the number of queries, keys and values, d_l and d_v are the dimension of query and key, the dimension of value, respectively. The output matrix of the Scaled Dot-Product attention layer is computed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_l}} \right) \mathbf{V} \quad (1)$$

Two encoded embeddings $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{C}}$ extracted from SBA blocks of words and characters are concatenated to generate the word-char embedding $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_U]$ where $U = W + C$ and $\mathbf{u}_u \in \mathbb{R}^{d_h}$. The word-char embedding \mathbf{U} is then fed into the word-char encoder block to generate the encoded word-char embeddings $\widehat{\mathbf{U}} = [\widehat{\mathbf{u}}_1, \widehat{\mathbf{u}}_2, \dots, \widehat{\mathbf{u}}_U]$. The word-char encoder block reuses the two stacked Bidirectional LSTM layers from the SBA block without using the attention layer.

Besides the lexical, syntactic and semantic information for each sentence extracted from the word and character branches, we consider the statistical information of individual sentence: the number of sentences in the same abstract, the index of sentence in the abstract, and the number of words in the sentence, which are represented by one-hot vectors $\mathbf{t}_1, \mathbf{t}_2, \mathbf{t}_3$. The statistical information equips each sentence with the ability to capture sequential and contextual properties related to other sentences within abstract. The statistical vectors are fed into a Multi-layer perception (MLP) to generate encoded

statistic embeddings $\widehat{\mathbf{T}} = [\widehat{\mathbf{t}}_1, \widehat{\mathbf{t}}_2, \dots, \widehat{\mathbf{t}}_T]$. The encoded statistic embeddings $\widehat{\mathbf{T}}$ are then concatenated with the encoded word-char embedding $\widehat{\mathbf{U}}$ to generate the word-char-stat embedding $\mathbf{Z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_Z]$ where $Z = U + T$ and $\mathbf{z}_z \in \mathbb{R}^{d_h}$. Again, one Bi-LSTM layer and one Dense layer are used to learn the sequence of word-char-stat embeddings, which combine both semantic, syntactic information of word-char encoded embedding and statistical information of statistic embedding, to generate the encoded word-char-stat embeddings $\widehat{\mathbf{Z}}$.

To further enhance the representation of sentences in term of language comprehension specifically in biomedical domains, we utilize BiomedBERT [16], which was pretrained on the PubMed corpus. The PubMed-based sentence embedding $\mathbf{P} = [\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_P]$ is concatenated with the encoded word-char-stat embedding $\widehat{\mathbf{Z}}$ before feeding into a Dense layer followed by a Softmax layer for classification. After training the Sen-Model at the sentence level, for each sentence, we extract the output of the Dense layer before the final Softmax layer and consider it as the final sentence-level embedding $\mathbf{s}_i \in \mathbb{R}^{d_L}$, where d_L is the dimension of a sentence-level embedding (i.e. the Softmax layer presents L outputs which match L labels of sentences in an abstract). The final sentence representation of the entire dataset is $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_S]$, where S is the number of sentences in the dataset and $\mathbf{s}_i \in \mathbb{R}^{d_L}$. The Sen-Model is optimized using Categorical Cross-Entropy:

$$L_{\text{Sen}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^L y_{ij} \log \widehat{y}_{ij} \quad (2)$$

where y_i, \widehat{y}_i and N are the true label, the predicted probability vector of sentence \mathbf{s}_i and the batch number, respectively. To examine the efficacy of the extracted sentence embeddings, we constructed two networks aimed at enhancing performance at higher levels by leveraging these embeddings. The networks are presented in the next subsections.

B. The regression model at the abstract level (Abs-Model)

Given the original dataset comprising of S sentences $[s_1, s_2, \dots, s_S]$, each sentence is now represented by a sentence-level embedding $\mathbf{s}_i, i = 1, 2, \dots, S$, extracted from the Sen-Model at the sentence level. To explore the sequential and contextual properties of sentences in one abstract, we group sentence embeddings in the same abstract to create the abstract representation $\mathbf{A} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_I]$ where $\mathbf{s}_i \in \mathbb{R}^{d_L}$ and I is the number of sentences in one abstract. The abstract representation \mathbf{A} is a sequence of sentence-level embeddings which is fed into the regression model at the abstract level (Abs-Model). The regression model at the abstract level (Abs-Model) is comprehensively presented at the left corner of the upper part of Fig. 1. The network includes two parts: Convolutional-Recurrent Neural Network (C-RNN) and Logistic Regression classifier.

Each abstract representation \mathbf{A} is now considered as a two-dimensional tensor which is fed into the convolution layers to extract essential features represented for neighbour sentences in one abstract. The two 2D-convolution layers in the C-RNN

TABLE I
MLP BASED NETWORK FOR CLASSIFICATION AT THE SEGMENT LEVEL

Blocks	Layers	Output Shape
F1	Dense (512) - Elu - BN - Dr(0.5)	512
F2	Dense (256) - Elu - BN - Dr(0.5)	256
F3	Dense (128) - Elu - BN - Dr(0.5)	128
F4	Dense (64) - Elu - BN - Dr(0.5)	64
F5	Dense (L) - Softmax	L

present similar settings in terms of kernel size, padding and the number of filters in each layer. Next, the Bi-RNN decoder is used for learning sequential relationship feature maps extracted from the convolutional layers. Finally, the Logistic Regression classifier receives the feature maps from the Bi-RNN decoder as input and generate predicted values $\hat{\mathbf{Y}}_{\text{abs}} = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_1]$ corresponding to the ground truth $\mathbf{Y}_{\text{abs}} = [y_1, y_2, \dots, y_1]$ where $\hat{y}_i, y_i \in \mathbb{R}^{d_L}$. Regarding the predicted value and the ground truth of one abstract, we form a predicted sequence $\hat{\mathbf{y}}_{\text{abs}} \in \mathbb{R}^{d_L \times I}$ and $\mathbf{y}_{\text{abs}} \in \mathbb{R}^{d_L \times I}$ by concatenating all the vectors of $\hat{\mathbf{Y}}_{\text{abs}}$ and \mathbf{Y}_{abs} , respectively. Then, the abstract-level Abs-Model is optimized using Binary Cross Entropy (BCE) loss on these predicted sequences, which can be written as:

$$L_{\text{Abs}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{d_L \times I} (y_{ij} \log(\hat{y}_{ij}) + (1 - y_{ij}) \log(1 - \hat{y}_{ij})) \quad (3)$$

where N is the batch number and the innermost sum presents the BCE loss for one abstract.

C. The classification model at the segment level (Seg-Model)

Given the extracted sentence embeddings, instead of generating all the segments with various lengths, we create fixed-length segments with the size of Q by grouping every Q consecutive sentences in one abstract. Each abstract of I sentences has $I - Q + 1$ segments. The i^{th} segment representation is described as $\mathbf{m}^{(i)} = [s_{Qi} s_{Q(i+1)} \dots s_{Q(i+Q-1)}]$, which is formed by concatenating Q continuous sentence embeddings. The corresponding label vector $\mathbf{y}_{\text{seq}}^{(i)}$ of the i^{th} segment is defined as:

$$\mathbf{y}_{\text{seq}}^{(i)} = \frac{\sum_{q=Qi}^{Qi+Q-1} \mathbf{y}_q}{\sum_{q=Qi}^{Qi+Q-1} \sum_{l=1}^L y_{ql}} \quad (4)$$

where $\sum_{l=1}^L y_{ql}$ is the sum of elements in the label vector \mathbf{y}_q of the sentence s_q . The fixed-length Q is set to 3 based on empirical experiments. The Seg-model at the segment level uses the same labels as which of the sentence level, meaning that all the sentences in a segment receive the label of that segment.

To classify segments, we use the MLP network which is shown in detail at table I. The network consists of five fully-connected blocks. The first four blocks present the same layers which perform Dense layer, ELU activation, Batch Normalization and Dropout, respectively. The output of the last block is used for segment-based classification task. Since the

TABLE II
DATASET STATISTICAL INFORMATION

Dataset	C	V	Train	Validation	Test
PubMed 20k	5	68k	15k/180k	2.5k/30k	2.5k/30k
PubMed 200k	5	331k	190k/2.2M	2.5k/29k	200/29k
NICTA-BIBOSO	6	17k	720/7.7k	80/0.9k	200/2.2k

labels of segment embeddings are no longer one-hot encoded, we use the Kullback-Leibler (KL) divergence loss for the segment-based classification task, which is defined as:

$$L_{\text{Seg}}(\theta) = \sum_{n=1}^N \mathbf{y}_{\text{seq}}^{(n)} \log \frac{\mathbf{y}_{\text{seq}}^{(n)}}{\hat{\mathbf{y}}_{\text{seq}}^{(n)}} + \frac{\lambda}{2} \|\theta\|_2^2 \quad (5)$$

where θ is the trainable parameters of the network, λ denotes the l_2 regularization coefficient experimentally set to 0.0001, N is the batch number, $\mathbf{y}_{\text{seq}}^{(n)}$ and $\hat{\mathbf{y}}_{\text{seq}}^{(n)}$ denote the ground-truth and the network output in a batch, respectively.

D. Inference with the entire system

Given the predicted labels of Abs-Model at the abstract level and Seg-Model at the segment level, referred to as $\hat{\mathbf{Y}}_{\text{abs}}$ and $\hat{\mathbf{Y}}_{\text{seg}}$, the final predicted labels of our proposed system is defined as:

$$\hat{\mathbf{Y}} = \lambda_{\text{abs}} \hat{\mathbf{Y}}_{\text{abs}} + \lambda_{\text{seg}} \hat{\mathbf{Y}}_{\text{seg}} \quad (6)$$

where λ_{abs} and λ_{seg} are the hyperparameters to control the predicted labels at the abstract level and the segment level.

III. EXPERIMENT AND RESULTS

A. Datasets

In this paper, we evaluate our proposed deep neural networks on two benchmark datasets: PubMed RCT [3] and NICTA-PIBOSO [4].

PubMed RCT: This dataset presents the largest and published dataset of text-based medical scientific abstracts. In particular, the PubMed dataset presents approximately 200,000 abstracts of randomized controlled trials. The total sentences in the PubMed dataset is around 2.3 million. Each sentence of each abstract is labeled with 'BACKGROUND', 'OBJECTIVE', 'METHOD', 'RESULT', or 'CONCLUSION' which matches its role in the abstract. The PubMed dataset proposed two sets of PubMed 20K and PubMed 200K, each of which presents three subsets of Training, Validation, and Test for training, validation and test processes, respectively.

NICTA-PIBOSO: This dataset is the official dataset of the ALTA 2012 Shared Task. The task was to build classifiers which automatically divide sentences to a pre-defined set of categories in the domain of Evidence Based Medicine (EBM), which are 'BACKGROUND', 'INTERVENTION', 'OUTCOME', 'POPULATION', 'STUDY DESIGN', 'OTHER'. Table II presents statistics information of these above datasets, where $|C|$ denotes the number of classes, $|V|$ denotes the vocabulary size. In the train, validation and test sets, we indicate the number of abstracts and the number of sentences separated by the slash (e.g. 15k/180k).

TABLE III

COMPARE OUR PROPOSED SYSTEMS WITH THE BASELINE ON THE TEST SET (F1 SCORE/PRESISION/RECALL)

Systems	PubMed 20K	NICTA-PIBOSO
bi-ANN [9] (baseline)	90.0/-/-	82.7/-/-
Sen-Model w/ word only	84.0/84.2/83.9	69.9/70.3/69.8
Sen-Model w/ word & char	84.2/84.2/84.2	70.0/70.3/69.8
Sen-Model w/ word & char & stat	89.5/89.7/89.3	77.9/77.9/77.9
Sen-Model w/ pre-trained sentence only	87.0/87.1/87.0	78.5/78.8/78.5
Sen-Model w/ sentence & word & char & stat	91.1/91.9/90.9	81.8/81.8/81.8
Abs-Model w/ word only	90.6/91.2/90.4	81.5/83.4/80.3
Abs-Model w/ word & char	90.7/91.3/90.5	81.4/83.1/80.5
Abs-Model w/ word & char & stat	91.5/91.8/91.2	81.2/82.6/80.1
Abs-Model w/ pre-trained sentence only	91.9/92.1/91.7	82.5/84.0/81.5
Abs-Model w/ sentence & word & char & stat	92.7/93.2/92.6	84.6/85.5/84.1

TABLE IV

COMPARE OUR PROPOSED SYSTEMS ON DIFFERENT LEVELS WITH THE BASELINE ON THE TEST SET OF TWO BENCHMARK DATASETS(F1 SCORE/PRESISION/RECALL)

Systems	PubMed 20K	NICTA-PIBOSO
bi-ANN [9] (baseline)	90.0/-/-	82.7/-/-
Sen-Model (Sentence)	91.1/91.9/90.9	81.8/81.8/81.8
Abs-Model (Abstract)	92.7/93.2/92.6	84.6/85.5/84.1
Seg-Model (Segment)	91.0/92.5/89.6	79.5/80.7/78.5
Combine-Model	92.8/93.4/92.7	85.3/86.5/84.5

B. Evaluation metric

In this paper, we follow the original paper [3], [4] which proposed the PubMed RCT and NICTA-PIBOSO datasets. We then use Precision, Recall, and F1 scores as evaluation metrics.

C. Experimental settings

We construct our proposed deep neural networks with the TensorFlow framework. While the deep neural network used for Sen-Model is trained for 30 epochs, we train Abs-Model and Seg-Model with 60 epochs. All deep neural networks in this paper are trained with the Titan RTX 24GB GPU. We use the Adam [17] method for the optimization. The learning rate for Sen-Model, Abs-Model and Seg-Model are 0.001, 0.003 and 0.001, respectively. A reduce learning rate scheme by a factor of 0.1 is set during training. The Bi-RNN decoder at Abs-Model uses Bi-LSTM for PudMed dataset and Bi-GRU for NICTA-PIBOSO dataset, respectively. The hyperparameters λ_{abs} and λ_{seg} are empirically set to 1 and 0.2, respectively. The two 2D-convolution layers in the C-RNN has the same padding with kernel size and number of filters set to (8, 3) and 16, respectively. The hidden states dimension d_h of all LSTM layers in Sen-Model is set to 128.

D. Experimental results

We first evaluate our proposed models at the sentence level with different input features: using only word sequence (Sen-Model w/ word only); using both word and character sequences (Sen-Model w/ word & char); using word, character, and statistics (Sen-Model w/ word & char & stat); using only sentence embeddings extracted from the pre-trained PudMed model (Sen-Model w/ sentence only); using all input features of word, character, statistics, sentence embeddings (Sen-Model w/ sentence & word & char & stat). The experimental results shown in Table III highlight that each input feature helps to further improve the performance of the Sen-Model at the

TABLE V

COMPARE OUR BEST MODEL WITH THE STATE-OF-THE-ART SYSTEMS ON TEST SET OF PUBMED 20K DATASET

Authors	Systems	F1-score
Yamada et al. [11]	Semi-Markov CRFs	93.1
Athur Brack et al. [18]	Transfer /Multi-task learning	93.0
Cohan et al. [19]	Pretrained BERT	92.9
Xichen Shang et al. [5]	SDLA	92.8
Jin and Szolovits [10]	HSLN	92.6
Gaihong Yu et al. [20]	MSM	91.2
Gongalves et al. [21]	CNN-GRU	91.0
Dernoncourt et al. [9]	bi-ANN	90.0
Agibetov et al. [22]	fastText	89.6
Our proposed model	BiLSTM-CRNN-MLP	92.8

TABLE VI

COMPARE OUR BEST MODEL WITH THE STATE-OF-THE-ART SYSTEMS ON THE TEST SET OF NICTA-PIBOSO DATASET

Authors	Systems	F1-score
Xichen Shang et al. [5]	SDLA	86.8
Athur Brack et al. [18]	Transfer /Multi-task learning	86.0
Yamada et al. [11]	Semi-Markov CRFs	84.4
Jin and Szolovits [10]	HSLN	84.3
Sarker et al. [23]	SVM	84.1
Cohan et al. [19]	Pretrained BERT	83.0
Dernoncourt et al. [9]	bi-ANN	82.7
M Lui [24], [25]	Feature stacking + Metalearner	82.0
Our proposed model	BiLSTM-CRNN-MLP	85.3

sentence level. The best performance at the sentence level is from the combination of all input features of word, character, statistics, sentence embeddings (Sen-Model w/ sentence & word & char & stat), presenting the F1/Precision/Recall scores of 91.1/91.9/90.9 and 81.8/81.8/81.8 on PubMed 20K and NICTA-PIBOSO datasets, respectively. This combination outperforms the model that uses only pre-trained sentence embeddings from the BERT model (Sen-Model w/ pre-trained sentence only), which scores 87.0, 87.1, and 87.0 on PubMed 20K, and 78.5, 78.8, and 78.5 on NICTA-PIBOSO. These results demonstrate the effectiveness of the proposed LSTM-based network in generating high-quality sentence representations. This is achieved by combining task-specific features based on words, characters, and statistical information within a specific context, along with pre-trained embeddings from the BERT model, which has a comprehensive understanding of medical domain language from large-scale medical corpora. In other words, the proposed LSTM-based network effectively integrates synergistic and diverse features, allowing model to consider both the overarching medical knowledge and the specific details of each sentence, resulting in superior sentence representations.

Leveraging sentence embeddings from the proposed Sen-model, the model at the abstract level (Abs-Model) further enhances the system performance. We achieve the F1/Precision/Recall scores of 92.7/93.2/92.6 on PubMed 20K and 84.6/85.5/84.1 on NICTA-PIBOSO datasets as shown in the lower part of Table III. The model at the segment level (Seg-Model), when being integrated into the system, shows efficiency in considering coherent dependencies of sentences in local regions within segments and recorrect sentences at the boundary of two label classes. The best system (Combine-Model), which combines of Abs-Model and Seg-Model, achieves the best result of 92.8/93.4/92.7 on PudMed

20K and 85.3/86.5/84.5 on NICTA-PIBOSO as shown in Table IV. This model also outperforms the baseline [9] by 1.0%, 2.8%, and 2.6% on PubMed 200K, PubMed 20K, and NICTA-PIBOSO datasets in terms of F1 scores, respectively.

Compared with the state-of-the-art systems as shown in Table V and Table VI, although our best system presents fundamental network architectures at the abstract level and the segment level when leveraging the well-presented sentence embeddings at the sentence level, we achieve very competitive results (top-4 on PubMed 20K and top-3 on NICTA-PIBOSO). This indicates that the role of the LSTM-based network (Sen-Model) at the sentence level is important to achieve comprehensive sentence representation, which can be effectively set as an initial foundation and leveraged in higher levels of segment level and abstract level to improve the model performance. Therefore, our future work is to investigate novel methods for further improving the model performance based on the well-presented sentence representation.

IV. CONCLUSION

This paper has presented a deep learning system for the Sequential Sentence Classification (SSC) task in medical scientific abstracts based on the motivation of improving sentence representation. By conducting extensive experiments, we achieved the best system that outperforms the baseline by 1.0%, 2.8%, and 2.6% on the benchmark datasets of PubMed 200K RCT, PubMed 20K RCT, and NICTA-PIBOSO regarding F1 scores, respectively. The results are highly competitive to the state-of-the-art systems on these two datasets. Particularly, our proposed LSTM-based network at the sentence level proves a vital role in generating comprehensive sentence representation, which can be served as a strong foundation for further exploring and improving the performance of the SSC task on higher contextual levels.

ACKNOWLEDGMENT

The work described in this paper is performed in the H2020 project STARLIGHT (“Sustainable Autonomy and Resilience for LEAs using AI against High Priority Threats”). This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 101021797.



REFERENCES

- [1] A. Rai *et al.*, “Query specific focused summarization of biomedical journal articles,” in *2021 16th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2021, pp. 91–100, doi:10.15439/2021F128.
- [2] H. S. Nguyen *et al.*, “Semantic explorative evaluation of document clustering algorithms,” in *2013 Federated Conference on Computer Science and Information Systems*, 2013, pp. 115–122.
- [3] F. Deroncourt and J. Y. Lee, “PubMed 200k RCT: a dataset for sequential sentence classification in medical abstracts,” in *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, Nov. 2017, pp. 308–313.
- [4] I. Amini, D. Martinez, and D. Molla, “Overview of the ALTA shared task,” in *Proceedings of the Australasian Language Technology Association Workshop 2012*, 2012, pp. 124–129.
- [5] X. Shang, Q. Ma, Z. Lin, J. Yan, and Z. Chen, “A span-based dynamic local attention model for sequential sentence classification,” in *Proc. ACL-IJCNLP*, 2021, pp. 198–203, doi: 10.18653/v1/2021.acl-short.26.
- [6] J. Lin, D. Karakos, D. Demner-Fushman, and S. Khudanpur, “Generative content models for structural analysis of medical abstracts,” in *Proceedings of the HLT-NAACL BioNLP Workshop on Linking Natural Language and Biology*, 2006, pp. 65–72.
- [7] P. Ruch, C. Boyer *et al.*, “Using argumentation to extract key sentences from biomedical abstracts,” *International Journal of Medical Informatics*, vol. 76, no. 2, pp. 195–200, 2007, doi: <https://doi.org/10.1016/j.ijmedinf.2006.05.002>.
- [8] S. N. Kim *et al.*, “Automatic classification of sentences to support evidence based medicine,” in *BMC bioinformatics*, vol. 12, no. 2, 2011, pp. 1–10, doi: <https://doi.org/10.1186/1471-2105-12-S2-S5>.
- [9] F. Deroncourt, J. Y. Lee, and P. Szolovits, “Neural networks for joint sentence classification in medical paper abstracts,” in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 2017, pp. 694–700.
- [10] D. Jin and P. Szolovits, “Hierarchical neural networks for sequential sentence classification in medical scientific abstracts,” in *Proc. EMNLP*, 2018, pp. 3100–3109, doi: 10.18653/v1/D18-1349.
- [11] K. Yamada, T. Hirao, R. Sasano, K. Takeda, and M. Nagata, “Sequential span classification with neural semi-Markov CRFs for biomedical abstracts,” in *Findings of the Association for Computational Linguistics: EMNLP*, 2020, pp. 871–877, doi: 10.18653/v1/2020.findings-emnlp.77.
- [12] Y. Peng, S. Yan, and Z. Lu, “Transfer learning in biomedical natural language processing: An evaluation of BERT and ELMo on ten benchmarking datasets,” in *Proceedings of the 18th BioNLP Workshop and Shared Task*, 2019, pp. 58–65, doi: 10.18653/v1/W19-5006.
- [13] Devlin *et al.*, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018, doi: 10.18653/v1/N19-1423.
- [14] J. Pennington, R. Socher, and C. Manning, “GloVe: Global vectors for word representation,” in *Proc. EMNLP*, 2014, pp. 1532–1543, doi: 10.3115/v1/D14-1162.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” 2023.
- [16] Y. Gu *et al.*, “Domain-specific language model pretraining for biomedical natural language processing,” *ACM Transactions on Computing for Healthcare (HEALTH)*, vol. 3, no. 1, pp. 1–23, 2021, doi: 10.1145/3458754.
- [17] P. K. Diederik and B. Jimmy, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2015.
- [18] A. Brack *et al.*, “Sequential sentence classification in research papers using cross-domain multi-task learning,” *International Journal on Digital Libraries*, pp. 1–24, 2024, doi: <https://doi.org/10.1007/s00799-023-00392-z>.
- [19] Cohan *et al.*, “Pretrained language models for sequential sentence classification,” in *Proc. EMNLP-IJCNLP*, 2019, pp. 3693–3699, doi: 10.18653/v1/D19-1383.
- [20] G. Yu, Z. Zhang, H. Liu, and L. Ding, “Masked sentence model based on bert for move recognition in medical scientific abstracts,” *Journal of Data and Information Science*, vol. 4, no. 4, pp. 42–55, 2019, doi: 10.2478/jdis-2019-0020.
- [21] S. Gonçalves, P. Cortez, and S. Moro, “A deep learning classifier for sentence classification in biomedical and computer science abstracts,” *Neural Comput. Appl.*, vol. 32, no. 11, p. 6793–6807, 2020, doi: 10.1007/s00521-019-04334-2.
- [22] A. Agibetov, K. Blagec, H. Xu, and M. Samwald, “Fast and scalable neural embedding models for biomedical sentence classification,” *BMC bioinformatics*, vol. 19, pp. 1–9, 2018, doi: <https://doi.org/10.1186/s12859-018-2496-4>.
- [23] A. Sarker, D. Mollá, and C. Paris, “An approach for automatic multi-label classification of medical sentences,” in *Proceedings of the 4th International Louhi Workshop on Health Document Text Mining and Information Analysis. Sydney, NSW, Australia*, 2013.
- [24] M. Lui, “Feature stacking for sentence classification in evidence-based medicine,” in *Proceedings of the Australasian Language Technology Association Workshop*, 2012, pp. 134–138.
- [25] D. Mollá, “Overview of the alta shared task: Piboso sentence classification, 10 years later,” in *Proceedings of the The 20th Annual Workshop of the Australasian Language Technology Association*, 2022, pp. 178–182.

Literature Books Recommender System using Collaborative Filtering and Multi-Source Reviews

Elena-Ruxandra Luțan
0000-0001-5363-9930

Department of Computers and Information Technology
University of Craiova, 200585, Craiova, Romania
Email: elena.ruxandra.lutan@gmail.com

Costin Bădică
0000-0001-8480-9867

Department of Computers and Information Technology
University of Craiova, 200585, Craiova, Romania
Email: costin.badica@edu.ucv.ro

Abstract—In this contribution, we present a method for obtaining literature books recommendations using collaborative filtering recommender system technique and emotions extracted from multi-source online reviews. We experimentally validated the proposed system using a book dataset and associated reviews that we collected from *Goodreads* and *Amazon* websites using our customized web scrapers. We show the benefits of using multi-source reviews by proposing a series of recommender system evaluation measures, which include single-source and multi-source recommendations similarity, recommendation algorithm usescases coverage and generated recommendations relevance.

I. INTRODUCTION

RECOMMENDER systems aim to help people choose different aspects of their life by relying on similar peers feedback. Traditional recommender system approaches of user ratings or browsing interaction can be enhanced by features extracted from user reviews, resulting in personalized recommendations results [2], [8].

In this paper, we extend our previous research on the topic of book recommendations considering emotions from social media book reviews ([5], [6]), by investigating the benefits of using multi-source reviews for creating an emotional categorization of literature books and provide book recommendations.

For analysis, we use a set of 1000 books rated on *Goodreads* website as "Best Books Ever", which we also used for our previous research [6]. Thus, we want to compare the results and show the improvements resulted by addition of the multi-source reviews.

For the 1000 books, we use two sets of book reviews that we collected from *Goodreads* and *Amazon* website, using our customized web scrapers.

We propose an user-based collaborative filtering recommendations algorithm, which identifies the top 5 most similar users to the user of interest based on the similarity of emotions present in their reviews. Then, a selection of 5 books enjoyed by the most similar users is recommended.

Lastly, we present three performance measures for evaluating our system: *Average Recommendations Similarity* which shows the average similarity between recommendations obtained using single-source reviews and multi-source reviews, *Algorithm Branch Coverage* which quantifies the recommendations algorithm branch used for providing recommendations

(i.e. information if the system used preferences of similar users to provide recommendations or suggested random book), *Relevance* which measures the relevance of the recommendations by evaluating the similarity between the reviewed book and recommended book.

For all three considered measures, we observe improvements in the recommendation process resulted from the addition of multi-source reviews information.

Compared to our previous works on literature books recommender systems, this paper includes: (1) comparison between usage of single-source and multi-source reviews for creating emotional book categorization, (2) introduction of a new book reviews dataset, collected from *Amazon* website using our customized we scraper, (3) presentation of recommendations quality performance measures and discussions revealing the benefits of using multi-source reviews.

The paper is structured as follows. In Section II, we present related works. Section III describes our proposed user-based collaborative filtering book recommendation algorithm using emotions from multi-source social media reviews. In Section IV, we provide and insight of the experimental dataset and discuss the experimental results. The last section presents our conclusions and future directions.

II. RELATED WORKS

Speciale et al. [10] implement and evaluate two book recommender systems using content-based and collaborative filtering techniques, using implicit user feedback. For experiments, two different data sources are used: a dataset regarding loans in all public libraries in Turin Italy and a dataset from Anobii social network. The authors acknowledge the benefits of using multi-source data, as it allows to include more users (beneficial for collaborative filtering) and to enrich book metadata (beneficial for content-based filtering).

Bouadjenek et al. [1] introduce a distributed collaborative filtering recommendation algorithm, which exploits and combines multiple data sources, aiming to improve the recommendations quality. The system is experimentally validated using two datasets, a bookmarking dataset and a movie dataset, using two training data ratings setups of 80% and 60% to predict the remaining ratings. Their experiments show the effectiveness

of the algorithm compared to state-of-the-art recommendation algorithms.

Liu et al. [3] propose a multi-source information approach to improve conversational recommender systems. For experiments they use two datasets with conversations centered around movies. Each movie is represented as an embedding built from the keywords identified in the reviews. During the interaction with the user, the system identifies user preferences based on dialog context, tags, entities, and predicts movies which might interest the user. In case user does not like the recommendation (e.g. user has already seen the movie), the system dynamically updates the knowledge about user preferences and provides new recommendations.

Roy and Ding [7] present a movie recommender system which uses different types of users feedbacks such as likes, comments, tweets, in addition to movie features (title, plot, genre director, actors). Their experiments show that the most accurate results are obtained when all feedback data is combined to represent the movie feature.

Schoinas and Tjortjis [9] propose a product recommendation system based on multi-source implicit feedback. The authors utilize different sources of information in addition to user purchase history, such as the the number of times users viewed an item and added it to the cart, in order to estimate the user preferences for items. The interaction score is computed using specific weights for each observation: viewing products has lowest weight, as it does only indicate that the user was interested to learn more about the product, while by adding it to cart or purchasing it, there are stronger indications that user preferred the product.

Toumy [11] discusses the idea of relying only on single most similar user when making recommendations, considering that by relying also on next most similar users it is likely to overwhelm the customer and lose credibility. Although this is an interesting idea, considering our limited number of users and reviews in the experimental dataset, we decided to use a small set of similar users.

III. SYSTEM DESIGN

We propose a method for obtaining valuable literature books recommendations using multi-source reviews and collaborative filtering recommendations technique. For analysis, we use a set of books, for which we collected reviews from *Goodreads* and *Amazon* websites using our customized web scrapers.

A book review is a tuple:

$$r = (book, user, date, stars, content)$$

where *book* represents the id which uniquely identifies a book, *user* represents the id which uniquely identifies an user, *date* represents the date when the review was written, *stars* represents a scaled rating provided by user - expressed as a natural number in interval [1, 5], *content* represents the content of the review.

At first, the book review shall be preprocessed in order to obtain the input for the recommender system. The input for the book recommender system is a tuple:

$$input = (book, user, emotions)$$

where *book* represents the id which uniquely identifies a book, *user* represents the id which uniquely identifies an user, and *emotions* is the frequency vector of emotions corresponding to the review content.

The emotions are extracted from the review content using the method we proposed in [4] which refers to applying standard NLP text preprocessing techniques (tokenization, lower casing, removal of stop words) to the review text, followed by a word-matching method of determining the emotions. We use an external file composed of adjectives and associated emotions. Following 35 emotions are considered: 'cheated', 'singled out', 'loved', 'attracted', 'sad', 'fearful', 'happy', 'angry', 'bored', 'esteemed', 'lustful', 'attached', 'independent', 'embarrassed', 'powerless', 'surprise', 'fearless', 'safe', 'adequate', 'belittled', 'hated', 'codependent', 'average', 'apathetic', 'obsessed', 'entitled', 'alone', 'focused', 'demoralized', 'derailed', 'anxious', 'ecstatic', 'free', 'lost', 'burdened'.

Then the collaborative filtering recommendation algorithm is applied (Algorithm 1).

Algorithm 1 User-Based Collaborative Filtering Recommendation Algorithm

```

1: Get user input review (book, user, emotions)
2: Create list of TOP 5 most similar users of user
3: Identify books enjoyed by similar users
4: if  $len(books) == nREC$  then
5:   Recommend books
6: else
7:   if  $len(books) > nREC$  then
8:     Recommend 5 random books from books
9:   else
10:    while  $len(books) < nREC$  do
11:      Add in books random book from the books dataset
12:    end while
13:    Recommend books
14:  end if
15: end if

```

The algorithm receives as input an user input review (*book, user, emotions*). Next, the top 5 most similar users are determined, by comparing the emotions of *user* with the emotions of all users who reviewed *book*.

Two users (A and B) are considered similar, if their emotion for *book* are matching above a given *threshold*. The similarity is computed using Cosine Similarity measure:

$$sim(A, B) = \frac{emotions_A \cdot emotions_B}{\|emotions_A\| \cdot \|emotions_B\|}$$

TABLE I
REVIEWS ENTITY DESCRIPTION

Field name	Field Description
Id	Integer number which uniquely identifies the review in our dataset
Book Id	The id which uniquely identifies the reviewed book
Book URL	The <i>Goodreads</i> or <i>Amazon</i> URL for the book, depending on the review source
Author Id	The id which uniquely identifies the user who wrote the review
Review Stars	The rating provided by the user together with the review - natural number in interval [0, 5]
Review Date	The date when the review was written
Review Content	The text content of the review

where $emotions_A$ and $emotions_B$ are the frequency vectors of emotions corresponding to reviews contents provided by users A and B.

The book preferences of the top 5 most similar users are determined and stored in an array *books*. A book preference is a book that was reviewed by a user with a rating of 4 or 5 stars.

At this stage, depending on the requested number of recommendations to be received ($nREC$), 3 use cases can be identified:

- 1) $len(books) = nREC$: in this case, all books from *books* are provided as recommendations.
- 2) $len(books) > nREC$: in this case, the array of similar users preferences contains more books than the requested number of recommendations to be received, which means that the system shall select $nREC$ books from *books* and provide them as recommendations. The $nREC$ books are randomly selected from the *books* array.
- 3) $len(books) < nREC$: in this case, the array of similar users preferences contains less books than the requested number of recommendations to be received, which means that the list of *books* shall be completed in order to contain $nREC$ books. The list is completed by selecting random books from the books dataset, which are not already available in the *books* array.

IV. EXPERIMENTS AND DISCUSSIONS

A. Dataset Overview

For analysis, we use a set of 1000 books rated on *Goodreads* website as "Best Books Ever", with the associated 129713 *Goodreads* reviews and 89401 *Amazon* reviews collected using our customized web scrapers.

For each review, we collected several parameters which are stored in a tabular file. The review entity with the collected parameters is presented in Table I.

On *Goodreads* website, we collected the first 6 reviews pages, with 30 reviews available per page, resulting in a maximum of 180 reviews per book, while on *Amazon* websites, we collected the first 10 reviews pages, with 10 reviews available per page, resulting in a maximum of 100 reviews per book.

TABLE II
REVIEWS DATASET STATISTICS

Statistic	Goodreads	Amazon
# of Collected Reviews	129713	89401
# of Reviews per book	[10, 180]	[1, 100]
Average # of Reviews per book	129	89
# of Reviews 5 stars	53305	54507
# of Reviews 4 stars	36547	22855
# of Reviews 3 stars	17559	7748
# of Reviews 2 stars	9572	2447
# of Reviews 1 stars	9183	1844
# of Reviews 'undefined' stars	3557	-




Fig. 1. Distribution of *Goodreads* Reviews Emotions

Each book has in average 129 *Goodreads* reviews (between 10 and 180) and 89 *Amazon* reviews (between 1 and 100) - Table II.

Based on user star ratings, 69% of *Goodreads* reviews are positive (4 or 5 stars), 14% are neutral (3 stars), 14% are negative (1 or 2 stars) and 3% are undefined, meaning that the user has only provided a review content, without assigning a scaled rating. For *Amazon* reviews, 86% are positive, 9% are neutral, 5% are negative.

In general, we observe that the majority of collected reviews are positive. When collecting the reviews, we used the default reviews sorting order, which refers to the fact that at first are displayed the reviews which obtained the most reactions from other users, in forms of likes or comments. This suggests that users tend to react rather on the positive reviews.

We observe a higher amount of emotions in *Goodreads* reviews compared to *Amazon* reviews (Figures 1 and 2). Although this was anticipated considering the number of reviews collected from each of the websites (1.5 times more




Fig. 2. Distribution of *Amazon* Reviews Emotions

Goodreads reviews than *Amazon* ones), the number of emotions identified in *Goodreads* reviews is 2.5 times greater than the number of emotions identified in *Amazon* reviews. This is justified by the fact that on dedicated books reviews websites (*Goodreads*), users tend to write more emotional elaborate reviews, compared to business oriented websites, such as *Amazon*.

B. Experimental Results and their discussion

We evaluated the performance our our system across 10 iterations. For each iteration, the reviews dataset was split as 80% training and 20% testing using stochastic sampling. Due to the fact that different number of reviews are available for each book (Table II), the training-testing split is done for each book reviews.

The training reviews dataset represents the reviews which are recorded in the system and are used for identifying the users which are similar to the user of interest. We have considered two types of training datasets:

- 1) *Single Source (SS)* which refers to usage of 80% *Goodreads* reviews for training
- 2) *Multi Source (MS)* which refers to usage of 80% *Goodreads* reviews and 100% *Amazon* reviews for training

The testing dataset remains the same in both cases (SS and MS) - 20% *Goodreads* reviews, and is used to simulate a population of users writing reviews and seeking for recommendations.

The training SS dataset contains 102854 reviews and training MS dataset contains 192255 reviews, while the testing dataset remains constant as 25955 reviews. This results in a total of 129775 recommendations being made with Recommendation Algorithm 1 in case of training SS and 129775 recommendations in case of training MS, as the system always recommends 5 books.

Let us define the parameters used for rigorous definition of our evaluation measures:

- *Recommendation space R* refers to the total number of possible recommendations, i.e. the total number of books available in the books dataset - in our case 1000.
- *User input space U* refers to the total number of user inputs u .

$$u = (\text{book}, \text{user}, \text{emotions}), \text{ where } \text{book} \in R$$

- *Test space T* refers to the subset of the input space $T \subset U$ used for experimental evaluation.
- A *recommendation* $f(u)$ refers to the output recommendations obtained when applying the recommendation algorithm 1. The output is a set of 5 books $r_i \in R$.

$$f: T \rightarrow R^5$$

$$f(u) = (r_1, r_2, r_3, r_4, r_5)$$

- *Recommendations similarity s* refers to the similarity between recommendations $f_1(u)$ and $f_2(u)$ provided for the same user input u using (1) only information available in

Goodreads reviews, respectively (2) information available in *Goodreads* and *Amazon* reviews for determining the recommendations.

$$s: R^5 \times R^5 \rightarrow [0, 1]$$

s is determined using Jaccard index.

$$s(f_1(u), f_2(u)) = \frac{|f_1(u) \cap f_2(u)|}{|f_1(u) \cup f_2(u)|}$$

- *Total number of relevant recommendations TNRR* refers to the amount of recommendations which are identified as relevant. A recommendation is considered relevant if the overall emotions of the reviewed book match the emotions of the recommended book $>$ *similarity_threshold*.

$$TNRR = \bigcup_{u \in T} (sim(u, f(u)) > \text{threshold})$$

$$sim(u, f(u)) \in [0, 1]$$

The similarity is computed using cosine similarity measure:

$$sim(u, f(u)) = \frac{\text{emotion}_{s_u} \cdot \text{emotion}_{f(u)}}{\|\text{emotion}_{s_u}\| \cdot \|\text{emotion}_{f(u)}\|}$$

We propose following performance measures for evaluating our recommendation algorithms:

- *Average Recommendations Similarity ARS* is the average of the similarity between recommendations provided using (1) only information available in *Goodreads* reviews, (2) information available in *Goodreads* and *Amazon*.

$$ARS = \frac{1}{|T|} \sum_{u \in T} s(f_1(u), f_2(u))$$

- *Algorithm Branch Coverage ABC* determines the type of recommendation that was provided, considering following types:

- 1) *General (GRL)*: Corresponds the case when the set of books enjoyed by similar users contains at least 5 books
- 2) *Random Fill (RF)*: Corresponds to the case when the set of books enjoyed by similar users contains books, but they are less than 5, which means that the list of recommendations has to be completed using a random range of books from the dataset
- 3) *Fully Random (FR)*: Corresponds to the case when the set of books enjoyed by similar users is empty, and the user is recommended a completely random list of 5 books from the dataset.

- *Relevance* shows what proportion of recommendations $f(u)$ are identified as relevant recommendations.

$$\text{Relevance} = \frac{TNRR}{(|R|)^5}$$

The ARS value obtained for all iterations are presented in Table III, which shows we obtained an average of 10.84%

TABLE III
AVERAGE RECOMMENDATIONS SIMILARITY MEASURE

Iteration	Identical Recomm.	ARS
I1	14232	10.96
I2	14118	10.80
I3	13989	10.70
I4	14086	10.78
I5	14011	10.71
I6	13966	10.68
I7	14354	10.98
I8	14254	10.90
I9	14214	10.87
I10	14429	11.03

TABLE IV
USER TYPE STATISTICS

Iteration	Reg. user	New user
I1	20534	5421
I2	20698	5257
I3	20741	5214
I4	20731	5224
I5	20637	5318
I6	20701	5254
I7	20727	5228
I8	20735	5220
I9	20741	5214
I10	20756	5199

similarity between the recommendations provided using as basis the SS, respectively the MS training reviews. This is a slightly higher value compared to our previous research [6] (2.43%) where we compared the similarity between the recommendations obtained using collaborative filtering and content based filtering recommendations techniques. The increase is due to the rather similar recommendation techniques used: collaborative filtering method using single-source and multi-source reviews for defining the similar users to the user of interest.

The 25955 testing reviews, were written by an average of 20700 registered users and 5255 new users. The average value is considered for the 10 iterations, and we observe small changes in the number of users in each category from one iteration to another - Table IV. This shows that, although we have a limited number of reviews, they are written by a quite high number of registered users, i.e. users who wrote a review belonging to the testing dataset and have written other reviews before, available in the training dataset.

We appreciate that this is an interesting result, as we have a very high number of users who wrote several reviews in the datasets. When we scraped the *Goodreads* and *Amazon* websites, we collected the top reviews ordered by their relevance and popularity, and, as a result, a high part of the top reviews is written by certain individual users.

The ABC measure (Table V) shows that when using MS training reviews, a higher number of recommendations are provided using the *General (GRL)* method compared to the SS training reviews, more precisely an average of 21254 compared to 20026 (computed for the 10 iterations). On the

TABLE V
ALGORITHM BRANCH MEASURE

Iter.	Training MS			Training SS		
	GRL	RF	FR	GRL	RF	FR
I1	21111	1125	3719	19919	2096	3940
I2	21196	1194	3742	19927	2176	4029
I3	21216	1174	3737	19964	2174	3989
I4	21396	1170	3563	20133	2175	3821
I5	21224	1156	3764	20025	2104	4015
I6	21250	1248	3634	19986	2213	3933
I7	21246	1176	3704	20013	2143	3970
I8	21251	1215	3676	20032	2192	3918
I9	21309	1206	3636	20143	2124	3884
I10	21339	1185	3631	20122	2116	3917

other hand, the average number of recommendations provided using random book choices decreases, *Random Fill (RF)* from 2151 (Training SS) to 1185 (Training MS) and *Fully Random (FR)* from 3942 (Training SS) to 3681 (Training MS). This is a result of the benefits of using multi-source reviews, as it leads to a higher amount of similar users and books enjoyed by the similar users, which means that more target users can receive proper personalized recommendations.

For the 25955 training reviews, a total of 199775 recommendations are provided using the proposed book recommendations algorithm. Table VI presents the average values of Relevance measures for the 10 iterations, depending on the *similarity threshold* value.

We observe that considering *similarity threshold* ≥ 0.5 , almost all provided recommendations are considered as relevant, as this rather low similarity value permits emotional vectors of reviewed book and recommended book to be quite different. High number of relevant recommendations is observed also for *thresholds* 0.6, 0.7, 0.8, while for 0.9, 54.11% of the recommendations are identified as relevant for Training MS and 49.30% for Training SS.

We remark that, regardless of the chosen *threshold* value, the number of relevant recommendations resulted when using the Training MS reviews is higher than the number of relevant recommendations resulted when using the Training SS reviews. This shows that using the information obtained from multi-source reviews leads to more accurate recommendations, relevant for the preferences of the user of interest.

V. CONCLUSIONS

In this contribution, we presented a method for making valuable book recommendations using emotion information extracted from single-source and multi-source social media in order to identify the similarities between books.

For experiments, we use a set of 1000 books and associated reviews collected from *Goodreads* and *Amazon* websites, using our customized web scrapers.

Our analysis shows the benefits of using book emotional information extracted from multi-source reviews, by considering a series of recommender system performance measures. We observed low similarity between recommendations obtained from single-source and multi source reviews (10.84%), which

TABLE VI
RELEVANCE MEASURE (AVERAGE VALUES PER 10 ITERATIONS)

Threshold	Training MS			Training SS		
	TNRR	not(TNRR)	Relevance	TNRR	not(TNRR)	Relevance
0.5	129729	46	0.9996	129719.9	55.1	0.9995
0.6	129629.6	145.4	0.9988	129524.6	250.4	0.9980
0.7	126462.9	3312.1	0.9744	126013.5	3761.5	0.9710
0.8	113468.7	16306.3	0.8743	111713.2	18061.8	0.8608
0.9	70230.4	59544.6	0.5411	63981.3	65793.7	0.4930

is justified by the fact that approx. 80% of the recommendations were generated using the *General* recommendations approach, i.e. generation of top 5 most similar users preferences and picking a random 5 books selections from their preferences, which leads to different 5 books recommendations.

Approx. 79% of users considered as seeking for recommendations are registered users, which means that the system knows their reading history and preferences expressed through previously written reviews, thus resulting in more accurate, personalized recommendations.

The *Relevance* measure provides an insight of the improvements in recommendations obtained through multi-source reviews, as for high similarity values, 5% more recommendations are identified as relevant.

As future work, we plan to improve the method for computing the similarities between users, by considering that two users are similar if their reviews show other common features such as keywords or tags, in addition to the similar emotions.

Another future direction is in regards of evaluating the received recommendations: the recommendations received by registered users could be compared with user reading history (i.e. the books which user has rated before), and the relevance of the recommendations can be evaluated based on the common features existing between the user's reviews, e.g. emotions, keywords, tags. Considering the fact that we have a quite high ratio of registered users vs. new users (4:1), we appreciate this would give interesting results about user readings preferences.

REFERENCES

- [1] M. R. Bouadjeneq, E. Pacitti, M. Servajean, F. Maseglia, and A. E. Abbadi. A distributed collaborative filtering algorithm using multiple data sources. In *The Tenth International Conference on Advances in Databases, Knowledge, and Data Applications*, 2018.
- [2] E. Hasan, M. Rahman, C. Ding, J. X. Huang, and S. Raza. Review-based recommender systems: A survey of approaches, challenges and future perspectives. *Computer Science*, 2405.05562, 2024.
- [3] H. Liu, Q. Cao, X. Huang, F. Liu, C. Zhang, and J. An. Multi-source information contrastive learning collaborative augmented conversational recommender systems. *Complex and Intelligent Systems*, 2024.
- [4] E.-R. Luțan and C. Bădică. Emotion-based literature book classification using online reviews. *Electronics*, 11(3412), 2022.
- [5] E.-R. Luțan and C. Bădică. Emotion-based literature books recommender systems. In *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, volume 35, pages 275–280, 2023.
- [6] E.-R. Luțan and C. Bădică. Experimenting emotion-based book recommender systems with social data. In *Information Technology for Management: Solving Social and Business Problems Through IT*, pages 164–182, 2024.
- [7] D. Roy and C. Ding. Multi-source based movie recommendation with ratings and the side information. *Social Network Analysis and Mining*, 11(76), 2021.
- [8] D. Roy and F. Shirazi. A review on multiple data source based recommendation systems. In *2021 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 1534–1539, 2021.
- [9] I. Schoinas and C. Tjortjis. Musif: A product recommendation system based on multisource implicit feedback. *15th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI)*, pages 660–672, 2019.
- [10] A. Speciale, G. Vallero, L. Vassio, and M. Mellia. Recommendation systems in libraries: an application with heterogeneous data sources. In *7th International workshop on Data Analytics solutions for Real-Life Applications*, 2023.
- [11] H. Toumy. Perfume project. 2019. https://hayatoumy.github.io/recommender_system/ [Accessed: (May 10, 2024)].

Toward a Framework for Determining Methods of Evaluation in Design Science Research

Julia Müller
University of Applied
Sciences Heilbronn,
Heilbronn, Germany
Email: jmueller8@stud.hs-
heilbronn.de
ORCID: 0009-0002-2184-
2357

Stefanie Würth
University of Applied
Sciences Heilbronn,
Heilbronn, Germany
Email: swuerth@stud.hs-
heilbronn.de
ORCID: 0009-0009-9469-
5213

Thomas Schäffer
University of Applied
Sciences Heilbronn,
Heilbronn, Germany
Email:
thomas.schaeffer@hs-
heilbronn.de
ORCID: 0000-0001-8097-
286X

Prof. Dr. Christian Leyh
Technical University of
Central Hesse (THM) –
University of Applied
Sciences, THM Business
School, Gießen, Germany
Email:
christian.leyh@w.thm.de
ORCID: 0000-0003-0535-
0336

Abstract—Evaluation is a key phase of design science research, particularly in design-oriented information systems, one that involves analyzing and solving problems to create artifacts. Because the nature of those artifacts varies based on the problem, they necessitate different methods of evaluation, and selecting an appropriate one first requires identifying appropriate criteria for evaluation. This paper aims to pinpoint those criteria by systematically reviewing the literature, with particular focus on identifying various criteria for evaluation, their frequency, their significance in evaluating artifacts, and their connection to specific methods of evaluation. The findings suggest a framework for choosing the most suitable methods of evaluation based on the defined criteria that can enhance the rigor and relevance of the evaluation phase in design science research.

Index Terms—Design Science Research, DSR, Methods of Evaluation, Criteria for Evaluation

I. INTRODUCTION

IN design-oriented information systems, design science research (DSR) is becoming increasingly important for solving specific problems and developing problem-solving artifacts [1], [2]. DSR is an especially suitable approach to creating artifacts that address a previously unsolved and important business problem and is therefore ideal for developing innovative solutions and prototypes in unexplored areas.

A central phase of DSR is evaluating developed artifacts to prove their usefulness, quality, and effectiveness [3]. Researchers using DSR face the challenge of selecting appropriate methods of evaluation [4]. For example, Ihlström Erikson et al. [5] have emphasized the need for detailed strategies in the process of evaluation that meet the specific needs of the project at hand. However, despite general agreement on the importance of evaluation in DSR, the literature lacks clear guidance on selecting evaluation-focused strategies and methods [4], [6]. The literature also shows inconsistencies within guidelines for selecting criteria and standards for evaluation [7]. Straßburg et al. [8] have additionally underscored the lack of guidance and the challenges of selecting an appropriate method for evaluating artifacts.

To provide a possible guideline for selecting appropriate methods of evaluation, the relationships between criteria and

methods for evaluation have to be determined first. Therefore, this paper examines the following research question:

Which evaluation criteria influence the choice of methods of evaluation in the context of design science research (DSR)?

The chief objective of our research project is to identify the relationships between artifacts, methods of evaluation, and criteria for evaluation based on a systematic literature review. By identifying those relationships, we aim to support decision-making about the selection and application of methods of evaluation in the context of DSR. To that end, criteria have been extracted from the literature, thoroughly analyzed, and summarized. With that summary, the paper provides a framework for selecting appropriate methods of evaluation that can support such decision-making in the selection of evaluation methods based on previously selected criteria in the context of DSR, all to improve the quality and effectiveness of evaluations.

In what follows, the paper first introduces relevant terms to establish a common understanding. Next, it presents the methodological approach used in the systematic literature review, after which it presents the results using the framework. Last, the results are summarized, and implications for future research are derived.

II. THEORETICAL BACKGROUND

This section presents the theoretical foundations and key concepts necessary for understanding our research question and the study guided by it. A firm understanding of those concepts provides a foundation for the analysis and discussion that follow. First, the approach of DSR and its relevance in information systems is explained. Second, the different methods of evaluation in the context of DSR are described in detail, along with their application and significance. Last, the specific criteria for evaluation that are important for evaluating artifacts developed in the context of DSR are discussed. Those principles make it possible to precisely address the research question and systematically justify the selection of appropriate methods of evaluation.

A. Evaluation in Design Science Research

Österle et al. [9] have classified research methods used in the context of information systems into *behavioral* and *design-oriented* approaches. According to both Hevner et al. [1] and Larsen et al. [2], DSR ranks among the design-oriented approaches. In information systems, DSR is understood as a problem-oriented paradigm aimed at creating innovative artifacts. Those artifacts, including models and methods, should enable the effective and efficient analysis, design, implementation, management, and use of information systems [1].

With roots in engineering, DSR typically entails phases such as *problem identification*, *requirements specification*, *design*, *evaluation*, and *communication* [3], [10]. Various approaches to evaluation differ in their applications depending on the type of artifact. Despite such diversity, the overarching goal of evaluation is to assess artifacts for their usefulness and effectiveness [11]. Evaluation is both a specific phase in the design–research cycle [1], [12] and an accompanying activity throughout the process [13]. The targeted assessment of intangible or tangible objects to justify their suitability is central to evaluation in DSR. According to Hevner et al. [1], evaluation is critical to the success of DSR.

B. Methods of Evaluation in DSR

Methods of evaluation are used to structure the evaluation and to transform the developed artifact into a research result [3]. Depending on the type of artifact, some methods of evaluation are more appropriate than others [3]. Different authors have developed different types of methods of evaluation and their application [7]. For example, Hevner et al. [1] divide the methods into five classes: *observational methods*, *analytical methods*, *experimental methods*, *test methods*, and *descriptive methods*. Peffers et al. [3], by contrast, distinguish only between *demonstration* and *evaluation*.

Peffers et al. [3], in discussing the variety of artifacts and methods of evaluation, have identified, among others, *logical reasoning*, *expert evaluations*, *technical experiments*, *subject-based experiments*, *action research*, *prototypes*, *case studies*, and *illustrative scenarios*. Some methods, including *demonstrations*, *literature reviews*, and *expert interviews*, are also commonly used but not explicitly defined in Peffers et al.'s [3], [11], [14] work.

C. Criteria for Evaluating Artifacts

March and Smith [10] have emphasized certain criteria when evaluating artifacts, including *construct*, *model*, *method*, and *instantiation*. *Constructs* are evaluated in terms of completeness, simplicity, elegance, comprehensibility, and ease of use [10], while *models* are evaluated according to their correspondence to real-world phenomena, completeness, level of detail, robustness, and internal consistency [10]. On the whole, operationality, efficiency, generality, and usability are important in evaluating *methods* [10], while *instantiations* are evaluated in terms of efficiency, effectiveness, and impact on the environment and users [10]. Although March and Smith [10] have answered the question of what to evaluate with different criteria, a specific method of evaluation has yet

to be assigned to the criteria, and an explicit assignment of criteria to methods of evaluation has not been investigated, either.

III. RESEARCH METHODOLOGY

We conducted a systematic literature review to identify published research addressing the selection of methods of evaluation and criteria in DSR. To ensure that our search for literature was comprehensive and structured, we followed vom Brocke et al.'s [15] approach, which we also supplemented with a forward and backward search as well as a visualization using a concept matrix according to Webster and Watson [16]. Webster and Watson [16] recommend using a concept-centered organizational framework in analyzing literature in order to enable the comparability of different contents within the literature. Vom Brocke et al. [15] have especially emphasized the strict documentation and presentation of the entire search process in scientific papers, for doing so ensures traceability for other researchers and optimal reusability.

Fig. 1 provides an overview of our methodology. The literature search was performed in the databases *AIS Electronic Library*, *IEEE Xplore Digital Library*, *Science Direct*, and *SpringerLink*. We conducted an electronic search of titles, abstracts, and keywords using the search terms [(“design science research” AND “DSR” AND “evaluation” AND “criteria”)] and [(“evaluation method” AND “criteria” AND “DSR”)]. The period from 2004 to 2023 was chosen in order to ensure that the results related to the DSR approach, with Hevner et al.'s [1] article, published in 2004, being the oldest included in our review.

The search yielded a total of 903 publications. After we excluded duplicates, the abstracts of 260 publications were reviewed for thematic relevance. After the abstracts, keywords and, if applicable, the full texts were analyzed, 24 publications relevant to our research's focus were identified. In addition to the database search, we performed a forward and backward search following the recommendations of Webster and Watson [16]. That final step raised the final number of publications to 30, all of which were read in detail and subsequently subjected to qualitative content analysis [17].

All publications were analyzed to identify methods of evaluation in relation to the evaluation criteria of the artifacts studied in the context of DSR. Six artifacts, eight methods of evaluation, and ten criteria for evaluation were identified in the 30 publications, all of which have already been analyzed in three works by March and Smith [10], Hevner et al. [1], and Peffers et al. [18]. The findings from the publications allowed extending the characteristics of evaluation, which were listed in a concept matrix next to the respective publications in order to derive correlations. In this paper, they are presented in a framework that serves as a guide for selecting methods of evaluation.

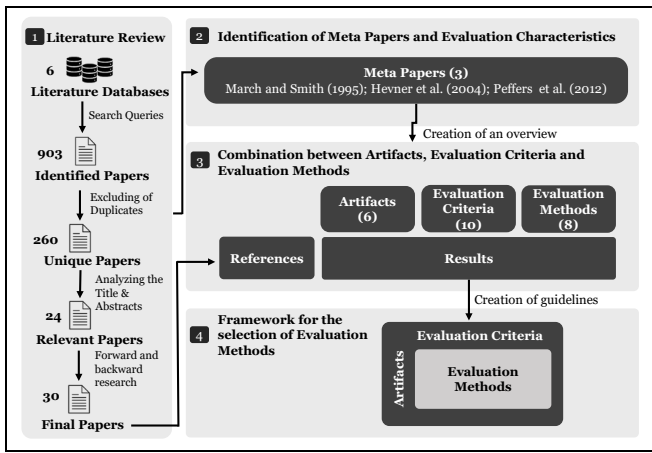


Fig. 1. Overview of the Research Methodology

IV. FINDINGS

In DSR, artifacts play a central role in developing and evaluating innovative solutions in information systems. Those artifacts, ranging from algorithms to frameworks to models, serve different purposes and require specific methods of evaluation in order to assess their effectiveness and applicability. Our literature review revealed six artifacts, ten criteria for evaluation, and eight methods of evaluation. Table I shows the results of the literature review in the form of a shortened concept matrix, as recommended by Webster and Watson [16]. The “SUM” column in Table I indicates the number of publications reviewed that refer to each specific concept.

In this section, the identified artifacts, criteria, and methods of evaluation are presented and explained in a structured manner, using the concept matrix as a visual aid. The developed framework is also presented, which serves as a decision-making aid for selecting methods of evaluation based on the previously selected criteria in the context of DSR.

The full concept matrix is available upon request from the authors.

A. Artifacts, Criteria, Methods of Evaluation, and their Relationships

Evaluation Criteria and their Application to Different Artifacts

Applicability: According to Baskerville et al. [12], *applicability* describes the ability of an artifact to achieve goals beyond its original purpose or to adapt to changing goals. The criterion of applicability should be used for evaluating frameworks and methods, for it allows a variety of methods of evaluation that in turn allow versatile application. The prototype was used most often for that purpose. For other artifacts, the criterion of applicability also allows evaluation with different methods, with the exception of the algorithm artifact.

Appropriateness: According to Akoka et al. [19], *appropriateness* is a criterion of evaluation used to ensure that the developed artifact is appropriate with respect to specific requirements and contexts. According to the definition of *appropriateness*, the methods need to be aligned with the intended outcomes. The criterion of appropriateness provides

the broadest range of applications for the framework, with nearly all methods being appropriate in the overall picture of the artifacts, with the exception of the literature review. However, for the artifact of the construct, no method is appropriate for the criterion.

TABLE I. CONCEPT MATRIX OF THE IDENTIFIED ARTIFACTS, CRITERIA, AND METHODS

		SUM
Artifact	Algorithm	3
	Construct	8
	Framework	9
	Instantiation	3
	Method	7
	Model	8
Criterion	Applicability	8
	Appropriateness	5
	Correctness	6
	Ease of Use	16
	Effectiveness	10
	Feasibility	3
	Functionality	4
	Understandability	4
	Usefulness	17
Validity	5	
Method	Case Study	10
	Demonstration Method	4
	Expert Evaluation	6
	Expert Interview	7
	Literature Review	3
	Prototype	7
	Subject-Based Experiment	3
	Technical Experiment	5
see references [19] – [48]		

Correctness: *Correctness* is defined as the degree of agreement between the results of the artifact and the expected results [9]. The criterion of correctness is covered by most methods of evaluation, with the exception of the expert interview and the literature review. There is no appropriate method of evaluation for the criterion of correctness in an instantiation.

Ease of Use: In the scientific literature, *ease of use* is described as the extent to which the artifact can be used effortlessly [49]. The expert interview is a common method of evaluating the usability of different artifacts. For algorithms, however, the search for suitable methods of evaluation to assess usability remains a challenge.

Effectiveness: In the context of methods of evaluation, *effectiveness* refers to the extent to which an artifact achieves its goal in a narrow sense, without considering situational aspects [50]. Although several methods of evaluation are suitable for most artifacts per the criterion of effectiveness, algorithms present unique challenges that limit the applicability of certain methods, including the case study.

Feasibility: According to Prat et al. [43], *feasibility* is a criterion of evaluation used to ensure that the developed artifact is viable in terms of technical, operational, and economic aspects. Per the definition of *feasibility*, the methods have to correspond with the outcomes of practical implementation and integration. The criterion offers a broad range of applications within the framework, with a variety of methods being suitable for evaluating feasibility, especially case studies, demonstrations, expert evaluations, and expert interviews.

Functionality: The *functionality* of an artifact refers to its ability to provide functions that meet predetermined and expected requirements when the artifact is used under specified conditions [51]. Functionality is often assessed in constructs, frameworks, and models, with expert interviews being a commonly used method. However, finding appropriate methods evaluating functionality in algorithms, instantiations, and methods can be more difficult.

Understandability: *Understandability* refers to the extent to which the artifact can be understood both at the global level and at the detailed level of elements and relationships within the artifact [9]. Understandability is typically evaluated in constructs, methods, and models. By contrast, evaluating understandability in algorithms, frameworks, and instantiations poses significant challenges, for appropriate methods are less likely to be identified.

Usefulness: According to Davis [49], *usefulness* is the extent to which a person believes that using a particular system would improve their work performance and is influenced by perception. Davis [49] has also defined *usefulness* as the extent to which an artifact positively influences an individual's task performance. Usefulness can be assessed for different artifacts, with frameworks and methods being particularly versatile. Case studies are commonly used, particularly for constructs, although identifying alternative methods can improve the assessment process.

Validity: According to Gregor and Hevner [52], *validity* means that the artifact works correctly—that is, achieves its goal in the right way. Validity is not appropriate for the artifacts of algorithms and instantiations because no appropriate methods have been identified. Methods could be analyzed for the remaining artifacts, with prototypes proving to be suitable for all four artifacts. The method and model artifacts, meanwhile, are particularly versatile because they use most of the methods.

Methods of Evaluation for Artifacts in DSR

Case Study: In a case study, the selected artifact is applied to a real situation. For that reason, the effects on the real situation can be evaluated [3].

Demonstration Method: By executing an artifact, the demonstration serves as a preliminary evaluation of the artifact's practical functionality [3].

Expert Evaluation: In an expert evaluation, an artifact is evaluated by one or more experts. Peffers et al. [3] give the example of Delphi studies.

Expert Interview: According to Sonnenberg and vom Brocke [14], expert interviews can be used repeatedly in all iterations of an evaluation. Although the questions within such interviews can be both objective and subjective, the answers are always subjective [11].

Literature Review: A literature review, primarily conducted in the early stages of an evaluation, serves to justify the problem and design objectives [13].

Prototype: A prototype is the implementation of the developed artifact and should serve to demonstrate its suitability and/or usefulness [3].

Subject-Based Experiment: The subject-based experiment method of evaluation is used to test the truth of a claim or hypothesis by using one or more test subjects [3].

Technical Experiment: A technical experiment is used to evaluate the technical performance of an implementation using real data, synthetic data, or no data, and the evaluation does not relate to performance in the real world [3].

B. Framework for Selecting Appropriate Methods of Evaluation

Based on the literature review and the concept matrix, we were able to gain important insights into the selection of methods of evaluation in the context of DSR. Through the structured presentation in the concept matrix, we were also able to determine which methods of evaluation are best suited for evaluating certain artifacts and criteria for evaluation.

To develop those mappings, we analyzed and evaluated the different combinations of artifacts, criteria, and methods of evaluation in an iterative process.

The resulting framework, shown in Table II, provides a systematic basis for selecting appropriate methods of evaluation, represented by numbers 1–8. Special attention was paid to the question of which methods are best suited to evaluate specific criteria. The specific scoring scheme is available upon request from the authors.

Depending on the artifact and the criterion, researchers can choose the appropriate methods to increase the quality and effectiveness of their evaluations. Shaded areas in the table indicate combinations for which no appropriate method of evaluation could be identified.

V. DISCUSSION

The results presented in this paper illustrate the complexity and multifaceted nature of selecting appropriate methods of evaluation in the context of DSR. In many of the reviewed publications, several methods are combined to evaluate an artifact, which contributes to the validation and increased significance of the results. For example, qualitative methods such as expert interviews can be complemented by technical experiments. Such multimethod approaches provide a more comprehensive view of the artifacts and strengthen the validity of the results of evaluation. The developed framework allows identifying versatile criteria for evaluation, including *effectiveness* and *usefulness*, that provide appropriate methods for each artifact.

TABLE II.
FRAMEWORK FOR SELECTING METHODS OF EVALUATION IN DSR

Evaluation Criteria \ Artifact	Applicability	Appropriateness	Correctness	Ease of Use	Effectiveness	Feasibility	Functionality	Understandability	Usefulness	Validity
Algorithm		2, 3	2, 3		1				1, 2, 3	
Construct	2, 3		1, 6, 7, 8	4, 6, 7, 8	3, 4, 6, 7, 8	2	2, 6	6, 7, 8	1	6, 8
Framework	1, 3, 4, 6	1, 4, 6	7	1, 3, 4, 6, 7	1, 4, 6		1, 4		1, 3, 4, 5, 6	1, 3, 6
Instantiation	7, 8	7, 8		4, 8	4, 6, 7	4, 5, 8			4, 5, 8	
Method	1, 2, 6, 7, 8	7, 8	3, 7, 8	1, 2, 3, 4, 6, 8	3, 4, 7, 8	3		1, 3, 4	1, 3, 4, 6	1, 3, 4, 6, 8
Model	4, 5, 6	1, 4	6, 8	1, 4, 8	2, 4, 5, 8		1, 4	1, 3, 4	1, 3, 4	1, 3, 4, 6, 8

Numbers are Methods of Evaluation: 1 - Case Study; 2 - Demonstration Method; 3 - Expert Evaluation; 4 - Expert Interview; 5 - Literature Review; 6 - Prototype; 7 - Subject-Based Experiment; 8 - Technical Experiment

The criteria of *applicability*, *appropriateness*, *correctness*, and *ease of use* proved to be particularly suitable. By contrast, *validity* with two and *functionality* as well as *understandability* with three unassigned artifacts were found to be less versatile.

A major limitation of our work is the assumption that all relevant criteria and methods were explicitly mentioned by the authors. That assumption could lead to bias, for established methods are cited more frequently, whereas newer approaches may be overlooked. The limited number of papers reviewed is another limitation. An extended search including additional databases and conference proceedings could improve the robustness of the results.

This paper does not provide a complete guide to evaluation but serves as a first step in selecting appropriate methods of evaluation based on the identified criteria. A detailed literature review in future studies could uncover more relevant methods and narrow gaps that currently prevent the development of a more comprehensive guide.

Future studies should additionally investigate the practical implications and challenges of the proposed methods through case studies. Doing so would validate the applicability and effectiveness of the methods in real-life scenarios and contribute to improving processes of evaluation.

VI. CONCLUSION

This paper provides important insights for researchers and their future research. The characteristics of evaluation identified in the concept matrix were mentioned several times in the different publications reviewed, which has allowed the creation of an initial guideline for selecting methods of evaluation. That guideline, based on the work of March and Smith [10], is intended to guide the selection of methods of evaluation based on the artifacts and criteria for evaluation.

We have answered our research question by identifying several correlations between methods of evaluation, artifacts, and criteria and by establishing a corresponding framework.

To date, no comparable study has considered all three dimensions—the artifact, the evaluation criterion, and the method of evaluation—in DSR. Our paper fills that gap and thereby provides a valuable foundation for future studies.

Overall, our paper demonstrates the need for a systematic, methodical approach to selecting methods of evaluation in DSR. Our paper additionally provides initial support for the selection of appropriate methods of evaluation and provides a basis for further research in the context of DSR. Although the presented framework is an important step, further research is needed. We recommend an extended literature review to identify additional evaluation criteria and their relationship to methods and artifacts. A more comprehensive review would narrow existing gaps in knowledge on the topic and develop a practical basis for a methodology for evaluation.

REFERENCES

- [1] A. Hevner, S. March, J. Park, and S. Ram, “Design Science in Information Systems,” *MIS Quarterly*, vol. 28, pp. 75-105, 2004.
- [2] Larsen, K. R., Lukyanenko, R., Mueller, R. M., Storey, V. C., VanderMeer, D., Parsons, J., and Hovorka, D. S., “Validity in Design Science Research,” in *Designing for Digital Transformation. Co-Creating Services with Citizens and Industry. DESRIST 2020*, Lecture Notes in Computer Science (LNCS), vol. 12388, S. Hofmann, O. Müller, and M. Rossi, Eds., Cham: Springer, 2020, pp 272–282.
- [3] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, “A Design Science Research Methodology for Information Systems Research,” *JMIS*, vol. 24, pp. 45-77, 2007.
- [4] A. Hevner, N. Prat, I. Comyn-Wattiau, and J. Akoka, “A pragmatic approach for identifying and managing design science research goals and evaluation criteria,” in *AIS SIGPRAG Pre-ICIS Workshop 2018*, San Francisco, USA, 2018.
- [5] C. Ihlstrom Eriksson, M. Åkesson, and K. Kautz, “Authentic and Concurrent Evaluation – Refining an Evaluation Approach in Design Science Research,” in *PACIS 2011 Proc.*, Brisbane, Australia, 2011.
- [6] E. Stoeckli, G. Neiditsch, F. Uebernickel, and W. Brenner, “Towards an Understanding of How and Why Design Science Research Scholars Evaluate,” in *Proceedings of ACIS 2017*, Hobart, Australia, 2017.
- [7] J. Venable, J. Pries-Heje, and Baskerville, R. “FEDS: A Framework for Evaluation in Design Science Research,” *Eur J Inf Syst*, vol. 25, pp. 77-89, 2016.
- [8] S. Straßburg, S. Kahlert, D. Stöffler, and T. Schäffer, “Identification of Issues in Design Science Research Evaluation – A Literature Review,” in *AMCIS 2021 Proceedings*, virtual conference, 2021.

- [9] H. Österle et al., "Memorandum on design-oriented information systems research," *Eur J Inf Syst*, vol. 20, pp. 7-10, 2011.
- [10] S. March, and G. Smith, "Design and natural science research on information technology," *Dec. Supp. Syst.*, vol. 15, pp. 251-266, 1995.
- [11] A. Cleven, P. Gubler, and K. M. Hüner, "Design alternatives for the evaluation of design science research artifacts," in *DESRIST 2009 Proceedings*, Philadelphia, USA, 2009.
- [12] R. Baskerville, M. Kaul, and V. Storey, "Genres of Inquiry in Design-Science Research: Justification and Evaluation of Knowledge Production," *MIS Quarterly*, vol. 39, pp. 541-564, 2015.
- [13] M. K. Sein, O. Henfridsson, S. Purao, M. Rossi, and R. Lindgren, "Action Design Research," *MIS Quarterly*, vol. 35, pp. 37-56, 2011.
- [14] C. Sonnenberg, and J. vom Brocke, "Evaluations in the Science of the Artificial – Reconsidering the Build-Evaluate Pattern in Design Science Research," in: *Design Science Research in Information Systems. Advances in Theory and Practice. DESRIST 2012*, Lecture Notes in Computer Science (LNCS), vol. 7286, K. Peffers, M. Rothenberger, and B. Kuechler, Eds., Berlin, Heidelberg: Springer, 2012, pp. 381-397.
- [15] J. vom Brocke et al., "Standing on the Shoulders of Giants: Challenges and Recommendations of Literature Search in Information Systems Research," *Comm. of the AIS*, vol. 37, pp. 205-224, 2015.
- [16] J. Webster, and R. Watson, "Analyzing the Past to Prepare for the Future: Writing a Literature Review," *MIS Quarterly*, vol. 26, pp. xiii-xxiii, 2002.
- [17] P. Mayring, "Qualitative Forschungsdesigns," in *Handbuch Qualitative Forschung in der Psychologie*, G. Mey, and K. Mruck, Eds., Wiesbaden: Springer, 2020, pp. 3-17.
- [18] K. Peffers, M. Rothenberger, T. Tuunanen, and R. Vaezi, "Design Science Research Evaluation," in *Design Science Research in Information Systems. Advances in Theory and Practice. DESRIST 2012*, Lecture Notes in Computer Science (LNCS), vol. 7286, K. Peffers, M. Rothenberger, and B. Kuechler, Eds., Berlin, Heidelberg: Springer, 2012, pp. 398-410.
- [19] J. Akoka, I. Comyn-Wattiau, N. Prat, and V. C. Storey, "Knowledge contributions in design science research: Paths of knowledge types," *Dec. Supp. Syst.*, vol. 166, Article 113898, 2023.
- [20] M. D. Ahmed, and D. Sundaram, "Design Science Research Methodology: An Artefact-Centric Creation and Evaluation Approach" in *ACIS 2011 Proceedings*, Sydney, Australia, 2011.
- [21] J. Barata, P. R. Da Cunha, and A. D. Figueiredo, "Self-reporting Limitations in Information Systems Design Science Research," *Bus Inf Syst Eng*, vol. 65, pp.143-160, 2023.
- [22] C. Basile, B. D. Sutter, D. Canavese, L. Regano, and B. Coppens, "Design, implementation, and automation of a risk management approach for man-at-the-end software protection," *Comput. Secur.*, vol. 132, Article 103321, 2023.
- [23] M. Bitzer, et al., "Managing the Inevitable – A Maturity Model to Establish Incident Response Management Capabilities," *Comput. Secur.*, vol. 125, Article 103050, 2023.
- [24] G. Bou Ghantous, and A.Q. Gill, A. Q., "Evaluating the DevOps Reference Architecture for Multi-Cloud IoT-Applications," *SN COMPUT. SCI.*, vol. 2, Article 123, 2021.
- [25] L. Bunnell, K-M. Osei-Bryson, and V. Y. Yoon, V. Y., "FinPath-light: Framework for a multiagent recommender system designed to increase consumer financial capability," *Dec. Supp. Syst.*, vol. 134, Article 113306, 2020.
- [26] B. M. Chaudhry, and J. Smith, "RefineMind: A Mobile App for People with Dementia and Their Caregivers," in: *The Next Wave of Sociotechnical Design. DESRIST 2021*, Lecture Notes in Computer Science (LNCS), vol. 12807, L. Chandra Kruse, S. Seidel, and G. I. Hausvik, Eds., Cham: Springer, 2021, pp 16-21.
- [27] Q. Deng, and S. Ji, "A Review of Design Science Research in Information Systems: Concept, Process, Outcome, and Evaluation," *PAJAIS*, vol. 10, 2018.
- [28] O F. El-Gayar, and B. D. Fritz, "A web-based multi-perspective decision support system for information security planning," *Dec. Supp. Syst.*, vol. 50, pp. 43-54, 2010.
- [29] D. A. Fischer et al., "Towards interactive event log forensics: Detecting and quantifying timestamp imperfections," *Information Systems*, vol. 109, Article 102039, 2022.
- [30] J. Forsberg, and T. Frantti, "Technical performance metrics of a security operations center," *Comput. Secur.*, vol. 135, Article 103529, 2023.
- [31] T. Matschak, F. Rampold, M. Hellmeier, C. Prinz, and S. Trang, "A Digitization Pipeline for Mixed-Typed Documents Using Machine Learning and Optical Character Recognition," in *The Transdisciplinary Reach of Design Science Research. DESRIST 2022*, Lecture Notes in Computer Science (LNCS), vol. 13229, A. Drechsler, A. Gerber, and A. Hevner, Eds., Cham: Springer, 2022, pp. 195-207.
- [32] S. Mdletshe, O. S. Motshweneng, M. Oliveira, and B. Twala, "Design science research application in medical radiation science education," *J. Med. Imaging Radiat. Sci.*, vol. 54, pp. 206-214, 2023.
- [33] A. L. Mesquida, and A. Mas, "Implementing information security best practices on software lifecycle processes: The ISO/IEC 15504 Security Extension," *Comput. Secur.*, vol. 48, pp. 19-34, 2015.
- [34] T. Mettler, M. Eurich, and R. Winter, "On the Use of Experiments in Design Science Research: A Proposition of an Evaluation Framework," *CAIS*, vol. 34, pp. 223-240, 2014.
- [35] B. Morschheuser, L. Hassan, K. Werder, and J. Hamari, "How to design gamification? A method for engineering gamified software," *Inform. and Soft. Tech.*, vol. 95, pp. 219-237, 2018.
- [36] J. A. Moutinho, G. Fernandes, and R. Rabechini, "Evaluation in design science: A framework to support project studies in the context of University Research Centres," *Evaluation and Program Planning*, vol. 102, Article 102366, 2024.
- [37] M. Muntean, R-D. Dănăiață, and L. Hurbean, "Applying Design Science Research for Developing Business Artifacts," *Proc. Comput. Sci.*, vol. 199, pp. 637-642, 2022.
- [38] T. Nagle, C. Doyle, I. M. Alhassan, and D. Sammon, "The Research Method We Need or Deserve? A Literature Review of the Design Science Research Landscape," *CAIS*, vol. 50, pp. 358-395, 2022.
- [39] K. Nahar, and A. Q. Gill, "Integrated identity and access management metamodel and pattern system for secure enterprise architecture," *Data & Knowledge Engineering*, vol. 140, Article 102038 2022.
- [40] F. K. de Oliveira, M. B. de Oliveira, A. S. Gomes and L. M. Queiros, "RECREIO: Floss as SAAS for sharing of educational resources," in: *Proc. of 12th Iberian Conference on Information Systems and Technologies (CISTI)*, Lisbon, Portugal, 2017.
- [41] M. Overeem, M. Mathijssen, and S. Jansen, "API-m-FAMM: A focus area maturity model for API management," *Inform. and Soft. Tech.*, vol. 147, Article 106890, 2022.
- [42] R. K. Pallasena, M. Sharma, and V. Krishnaswamy, "A Study of Interaction, Visual Canvas, and Immersion in AR Design: A DSR Approach," *AIS Trans. on HCI*, vol. 14, pp. 390-425, 2022.
- [43] N. Prat, I. Comyn-Wattiau, and J. Akoka, "A Taxonomy of Evaluation Methods for Information Systems Artifacts," *JMIS*, vol. 32, pp.229-267, 2015.
- [44] I. G. A. Premananda, A. Tjahyanto, and A. Mukhlason, "Design Science Research Methodology and Its Application to Developing a New Timetabling Algorithm, in: *Proc. of IEEE CyberneticsCom 2022*, Malang, Indonesia, 2022.
- [45] J. Pries-Heje, J. Venable, and R. L. Baskerville, "RMF4DSR: A Risk Management Framework for Design Science Research," *Scand. J. of Inform. Syst.*, vol. 26, pp. 57-82, 2014.
- [46] S. Pulparambil, Y. Baghdadi, and C. Salinesi, "A methodical framework for service-oriented architecture adoption," *Inform. and Soft. Tech.*, vol. 132, Article 106487, 2021.
- [47] A. Shrestha, A. Cater-Steel, M. Toleman, and T. Rout, "The Role of International Standards to Corroborate Artefact Development and Evaluation," in: *Software Process Improvement and Capability Determination. SPICE 2017*, Communications in Computer and Information Science, vol. 770, A. Mas, A. Mesquida, R. O'Connor, T. Rout, and A. Dorling, Eds., Cham: Springer, 2017, pp. 438-451.
- [48] P. M. Wiegmann, M. Talmir, and S. B. De Nijs, "Forging a sharper blade: A design science research approach for transition studies," *Environ. Innov. and Soc. Tran.*, vol. 48, Article 100760, 2023.
- [49] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quarterly*, vol. 13, pp. 319-340, 1989.
- [50] P. Checkland, and J. Scholes, *Soft systems methodology in action*. Hoboken, New Jersey, USA: John Wiley & Sons.
- [51] ISO/IEC/IEEE, *ISO/IEC/IEEE 24765:2017 Systems and software engineering — Vocabulary*, Edition 2, September 2017.
- [52] S. Gregor, and A. R. Hevner, "Positioning and presenting design science research for maximum impact," *MIS Quarterly*, vol. 37, pp. 337-355, 2013.

Pathomorphological Diagnosis Process Modeling for Machine Learning Algorithms' Applying

Małgorzata Pańkowska,
Mariusz Żytniewski
0000-0001-8660-606X
0000-0003-2170-1191
University of Economics
in Katowice, 1 Maja St.
40-287 Katowice, Poland,
Email: {pank,
zyto}@ue.katowice.pl

Mateusz Kozak,
Krzysztof Tomaszek
0009-0003-5978-1745
0009-0004-7272-1891
iMedLab, Pilsudskiego St.
41-902 Bytom, Poland
Email: {mkozak,
ktomaszek}@imedlab.pl

Dominik Spinczyk
0000-0003-0068-2948
Faculty of Biomedical
Engineering, Silesian University
of Technology, Roosevelta St.,
44-10 Gliwice, Poland
Email:
dominik.spinczyk@polsl.pl

□ **Abstract**— Business process management is oriented towards improving processes to best support people, who are working in them. Recent innovations in the area of artificial intelligence (AI), machine learning (ML), Internet of Things (IoT), and distributed systems have provided opportunities for new technologies applications, including process automation. This paper aims at the pathomorphological diagnosis (PD) process modeling for the ML solution implementation. The research method covers the PD laboratory case study. Authors argue that the PD process requires detailed analysis for its digitalization, automation, and combining with ML applications. Authors presented PD process models in BPMN notation, including laboratory equipment and emphasizing data and ML algorithms which are to be utilized in PD process digitalization for appropriate diagnosis for patients. Authors have found and emphasized that implementation of ML/AI algorithms is strongly based on fundamental process modeling.

Index Terms— Pathomorphology Diagnosis, Digital Image Analysis, Process Modeling, BPMN, Exploratory Data Analysis, Machine Learning.

I. INTRODUCTION

Process management has been developed for the last thirty years. In 2022, van de Aalst [27] has widely recommended the process mining concept, which has changed the process thinking way. Maas et al. [15] has done research on healthcare processes and they have developed automated medical reporting. Earlier, the process engineering focused on business or production processes' design and implementation. Authorities, i.e., Marlon et al. [16] argue that process management is a comprehensive body of tools, techniques, methods, and entire methodologies to support all stages of the business process lifecycle. They emphasize that the process approach is popular in industrial engineering, operations management, quality management, human resource management, corporate governance, computer science, and information systems engineering. However, this study aims to model the pathomorphology diagnosis (PD) process for generating valuable data

for ML and AI. In this study, researchers propose to integrate process knowledge from two disciplines, i.e., science of quality and management, and science of pathomorphology. The authors claim that only through combining managerial knowledge and medical knowledge, the pathomorphological diagnosis process can be effectively and efficiently controlled and surveilled [2].

Researchers in this study have noticed that there is lack of research on the PD ecosystem architecture development to ensure the PD process surveillance and control. They are able to develop the system architecture as a setting for the process realization and management. Authors have noticed that so far, the PD processes at various laboratories were utilized for human decision-making. This study presents the PD processes, which are to provide high quality pre-processed data for ML and AI-based analysis and prediction. Various researchers of ML and AI mostly focus just on reasoning and prediction. However, in this study, authors emphasize the issue of data provenance and governance, as well as the PD process modeling. Authors know that analysis of histopathological images is a highly required standard for patient therapies, hence this study results are expected to be valuable for the PD laboratory managers, pathomorphologists, and researchers of machine learning. Authors agree with Kratsch et al. [14] that process management is data-driven as well as model-driven approach. Data-driven orientation focuses on data for process discovery, analysis, monitoring, and process mining. It is true that data-driven approaches help analyze and monitor processes and to determine how efficient and effective their performance is. However, data concerns task attributes as well as tasks' results. In this case study, authors propose the ML and AI algorithms for analysis of the digitalized PD process results. The authors conducted a survey of literature [12, 20] in the field of pathomorphology diagnosis processes, and they searched for publications on PD processes in repositories, i.e., PubMed, Web of Science (WoS), and Scopus. Research works pub-

□□□□ This work was not supported by any organization

lished in 2000-2004 and in 2019-2023 were examined, revealing noticeable changes in popularity of the searched terms. Particularly, the researchers have noticed a decrease in the occurrence of terms such as 'results' and 'methods' and an increase of popularity of the term 'patient'. This indicates a shift in the articles' focus towards patient-oriented perspectives. Hence, in this study authors focused on modeling of patient-oriented process.

This paper is structured in the following ways. Section 2 is on the PD process modeling for ML and AI algorithms' implementation. The last section covers conclusions and recommendations.

II. PATHOMORPHOLOGICAL DIAGNOSIS PROCESS DECOMPOSITION

The aim of this paper is to provide a guide on how to select the best strategy for modeling the pathomorphological diagnosis process. Authors considered theoretical and practical perspectives. They began from the literature survey. Authors have reviewed the following repositories: Scopus, Association for Information Systems electronic Library (AISeLib), PubMed, IEEEExplore, and Sage Journals. The search query used to search titles and abstracts was: "process modeling" AND "machine learning". The search included articles written in English and conference proceedings between January 2013 and January 2024. The inclusion criteria were used to select the publication first based on their title and then on their keywords and abstracts. Extra sources were added from the references of the selected publications. Although, in general, 1227 publications have been found, just 21 publications were finally selected. However, next, 14 publications have been excluded, and finally only 8 articles (i.e., [5, 7, 9, 14, 18, 24, 25, 28]) were considered as valuable contributions on the process modeling for ML/AI application.

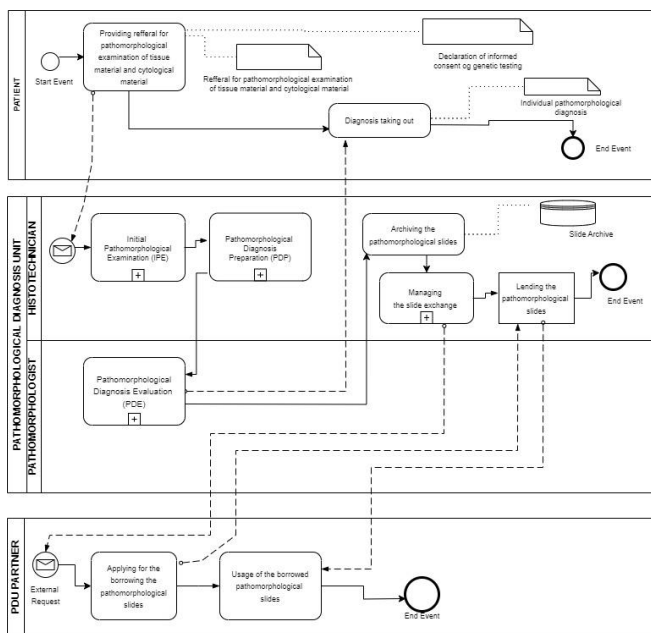


Fig 1. Pathomorphological Diagnosis Unit (PDU) main process

The others are mostly focused on ML and AI applications, but process modeling is not precisely presented. The other source of knowledge was regulations and guidelines of the Ministry of Health [17, 19, 22]. However, the individual in-depth interviews in the pathomorphological diagnosis unit (PDU) were the most valuable research method suitable for the identification of the PDU processes. According to the authors the pathomorphological diagnosis unit (PDU) process can be defined as a set of sequential and parallel activities, executed by medical and technical team members with different competencies and capabilities (see Fig. 1).

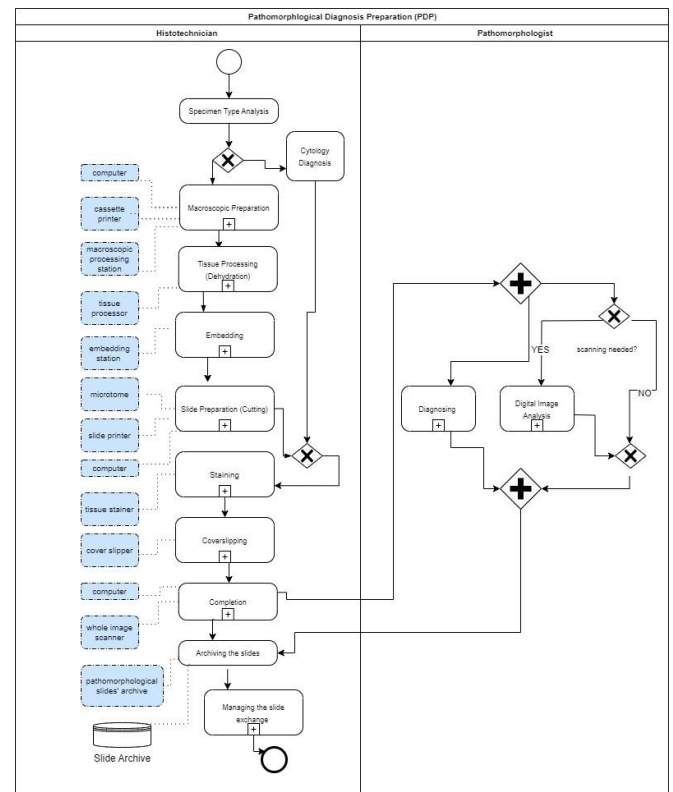


Fig 2. Pathomorphological Diagnosis Preparation (PDP) subprocess

The PDU process requires usage of special equipment and tools with the goal of high-quality treatment of materials provided to diagnose. Taking into account the literature survey, other authors focus on the material specification and its analysis. Generally, various modeling strategies are applied and characterized by their own granularity levels, data acquisition methods, modeling techniques and notations. In this paper, authors present modeling the main PDU process (Fig. 1.) and its decomposition into three subprocesses, i.e., Initial Pathomorphological Examination (IPE), Pathomorphological Diagnosis Preparation (PDP), and Pathomorphological Diagnosis Evaluation (PDE).

Generally, the pathomorphological diagnosis is a complex set of tasks, hence the process network and hierarchy of processes are necessary to order the PDU activities. The PDU has two main actors, i.e., histotechnician and pathomorphologist. The ordering of their tasks permits segregation of duties, and

finally, the pathomorphologist should focus on image recognition and evaluation, while the histotechnician should be responsible for the slides' preparation. In Fig. 1 and in Fig. 3 tasks are combined with data objects, which are generated in those tasks and which are further processed through the ML algorithms.

In Fig. 2, the PDP process tasks are combined with the equipment, which is installed in the PDU laboratory. Although the BPMN notation does not include specification of equipment in a process model [4], authors want to emphasize that equipment is essential in the PDU process. At first, after receiving the specimen, the histotechnician answers the question if it is a cytology tissue or other histology material or cytoblocks. The cytology material is analyzed in a separate way in the Cytology Diagnosis task. The Macroscopic Preparation task is performed by a histotechnician, who is using the macroscopic processing station, the cassette printer, and computer for recording the histopathology material data in the database. Next, the histotechnician performs the Tissue Processing task and uses the tissue processor. The Embedding task is also performed by the histotechnician at the embedding station. Next, the histotechnician performs the following tasks: the Slide Preparation with the microtome, slide printer, and computer for data recording, the Staining task with the tissue stainer, the Coverslipping task with the cover slipper, and the Completion task with the whole image scanner and computer for recording the scan description.

The results of histotechnician work are transferred to the pathomorphologist for diagnosis and digital image analysis. Finally, the pathomorphological slides are archived for further studying. The business process model and notation (BPMN) process modeling allows to clearly visualize workflows, making easy-to-interpret tasks and present relations between extensive sets of actions and decisions. Actually, the PDP process tasks are realized manually, but soon, the AI solutions and the robots' implementation will enable the PDP process digitalization and full automation. The PDP sub-process model is to be further used in cost effectiveness analysis of PDU laboratories.

Authors of this study assume that process modeling starts with the process discovery phase. Through the in-depth interviews, authors decide to focus on producing a detailed description of the business process as it currently exists (i.e., "as-is" process). Further, during the process analysis, application of analytical tools and techniques is to allow determining of process current weaknesses and ambiguities. The process redesign addresses the most important weaknesses and provides the "to-be" process model. That model can be used as the basis for process digitalization and automation. The PDP subprocess is also complex and as such requires further decomposition, which is presented in Fig. 2. Taking into account the literature survey presented in the section 2 of this paper, authors have focused on the Digital Image Analysis subprocess, which is the most important in the whole work of a PDU. The Digital Image Analysis (DIA) process model is included in Fig. 3. That process requires a strong involvement

of histotechnician and pathomorphologist in tasks for digital images' preparation. Defectively performed tasks require repetitions, which are costly and time consuming.

In the Digital Image Analysis process, the pathomorphologist performed tasks, which provide data needed for diagnoses based on Machine Learning (ML) as well as Artificial Intelligence (AI) reasoning. The process in Fig. 3 includes specifications of data objects that are necessary for ML/AI analyses. The tasks identified in the DIA process are further characterized in Table 1.

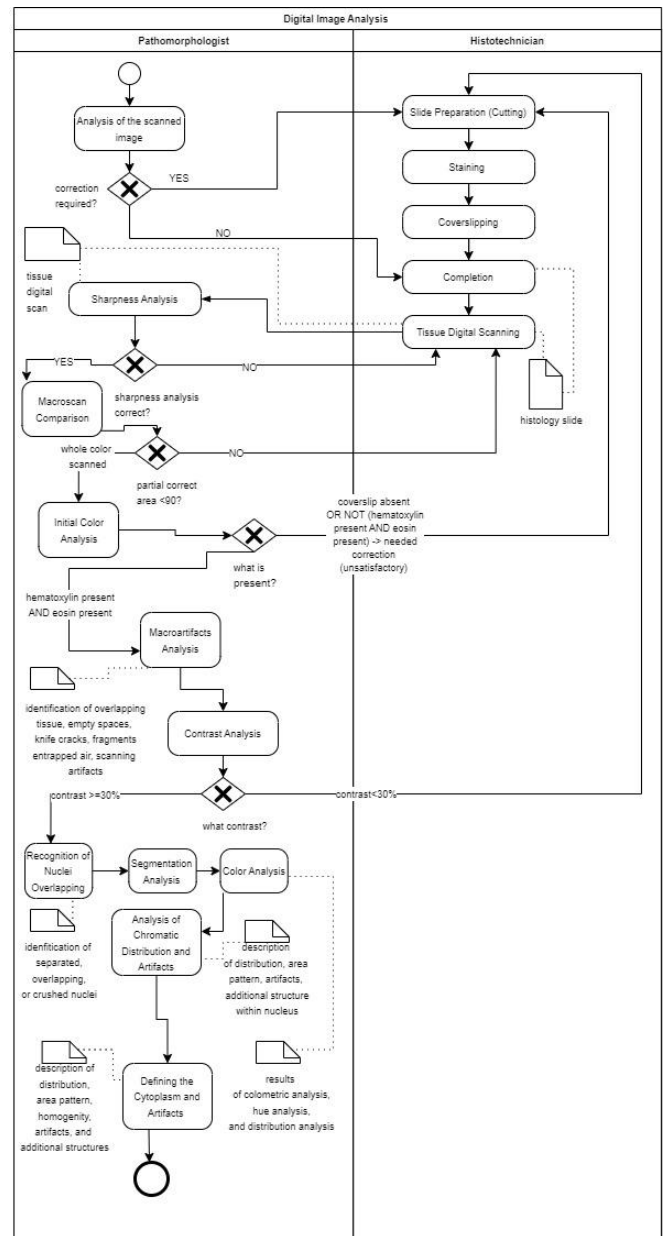


Fig 3. Digital Image Analysis (DIA) subprocess

The tasks indicated in the Tables I-II and in the Fig. 1 are performed manually by the pathomorphologist for the processing of the digital image of the specimen. By specifying and characterizing them in a process model, it is possible to

design an image processing workflow that will automate the Digital Image Analysis process. This workflow will consist of three stages:

1. Each task identifies an existing computer vision algorithm that can be used to automatically process the original image into a feature map that is appropriate for evaluating the image in the job.
2. The resulting feature map, in point 1, and the possible error classes in a given task are used to train artificial intelligence models, e.g., a deep convolutional network for each task, which classifies the image as correct or finds an error class [13, 26]. In the case of

transforming the image to a different form (e.g., histogram), an expert system based on detailed expert decision-making rules will be applied [1].

3. On the basis of the classification assigned to the image in a given task, in point 2, a decision is made to further process the image of the original specimen in accordance with the process diagram.

According to the literature survey, there are researchers who undertake the challenge to combine process modeling and ML/AI solution implementation [7, 8, 9, 14, 18, 24, 25, 28]. They argue that machine learning (ML) models offer different and wide-ranging capabilities to improve business processes.

TABLE I.
DESCRIPTION OF THE TASKS OF THE DIGITAL IMAGE PROCESSING PROCESS OF THE SPECIMEN

No	Task name	Purpose of the task	What is the algorithm used/ what happens to the image	Evaluation criterion	Indication of the reasons of irregularities/ linking the result of the task with other elements of the process
1	Sharpness Analysis	Assessing whether the specimen is scanned sharply over the entire surface	You can apply an algorithm to detect the edges [23] in the image and evaluate their thickness. If the thickness is too thick, the user has the impression of blurring.	There must be sharpness on the surface of the preparation, e.g.,: $\geq 80\%$	It results directly from the process of scanning the specimen in a histopathology scanner.
2	Macroscan comparison	Comparison of whether the surface of the scanned specimen in the micro view corresponds to the surface in the macro view	An algorithm can be applied to capture images in micro and macro view. Image registration can be based on elements of ML/Artificial Intelligence [3, 21]. If these views do not coincide, the content of the micro offense is distorted.	Images must overlap at least, e.g., 80%	Abnormal tissue detection by histopathology scanner algorithm.
3	Initial Color Analysis	Color separation; Is there hematoxylin? Is there eosin? Are there other colors?	The proportion of the number of pixels in the image is checked, and compared with colors in the shade range for hematoxylin and eosin. The image histogram processing algorithm can be applied [6,10].	1. Is there a correct number of pixels with color in the hematoxylin range? 2. Is there a correct number of pixels with color in the range? Is there eosin? 3. Are there other colors in the image histogram and what is their distribution?	Lack of hematoxylin, eosin - failures in staining task. In this case, the slide is discarded. If everything is monochrome, it means that there is no cover slip (the specimen is not sealed) - we reject the specimen from further analysis. In this case, the return to cover slipping is necessary. Recording that the process failed at this stage.
4	Macro Artifacts Analysis	Analysis of the surface of the preparation for the occurrence of specific phenomena.	A local correlation detection algorithm can be used [6, 10]. 1. Do the shapes of the "white" field form repetitive fields, e.g., the line. Do the shapes of the "white" field coincide with the lack of a pattern?	The following abnormal patterns can be detected: 1. Qualitative - slots 2. Mesh/disintegrating - poorly fixed 3. The tissue overlaps geometrically (triangle, square) with a more saturated area.	1. Damaged knife when slicing slides in a microtome job. 2. Error in tissue processing. 3. Poorly applied material to the slide.
5	Contrast Analysis	Is the scan image contrasting, e.g., between hematoxylin or eosin?	Algorithm for detecting the directional ratio of simple contrast [6,10]. 1. Edges are low contrast. 2. Too big contrast	The contrast must be within the given range.	1. Improper dyeing. 2. Improper fixation. 3. Dried material - tissue processor.

TABLE II.
DESCRIPTION OF THE TASKS OF THE DIGITAL IMAGE PROCESSING PROCESS OF THE SPECIMEN- CONTINUED

No	Task name	Purpose of the task	What is the algorithm used/ what happens to the image	Evaluation criterion	Indication of the reasons of irregularities/ linking the result of the task with other elements of the process
6	Recognition of Nuclei Overlapping	Color separation to hematoxylin (violet-blue color). Determine the cell nuclei, do they have a lot of outlines?	Segmentation of cell nuclei by dynamic thresholding based on the image histogram [6, 10], corresponding to the color of the hemotoxiline (violet-blue color). - there is a separation of pixels lying in this range: analysis of envelopes, shapes.	A high degree of overlapping of the testicles indicates a coarse cut.	In the case as it is written on the left, the material is too coarsely cut.
7	Segmentation Analysis	Segmentation of nuclei and cell bodies in the image	Segmentation algorithms corresponding to tiny numerous structures in pathological images. The possibility of using AI elements: deep neural networks in two modes of processing image fragments: – dividing the image into small areas and combining them [11]; Whole Slice Image Analysis [13, 29].		
8	Color Analysis	Analysis of the shade of dyes in the context of the thickness of the section, uniformity of coloration, shade change, color zoning	Analysis of color distribution in image chromatogram [3].	The base shade should be uniform within the cutter	Staining, reagent consumption, pH disturbances, poor dewaxing
9	Analysis of Chromatin Distribution and Artifacts	Analysis of chromatin distribution (haematoxylin) within the nuclei, surface analysis	Analysis of color distribution in the image chromatogram, analysis in the range of chromatin (chemotoxylin) colors [6, 10]		Distribution, area pattern, artifacts, additional structures within nucleus
10	Defining the Cytoplasm and Artifacts	Analysis of color within cell bodies	Analysis of color distribution using segmentation results from task 7 [6, 10].		Distribution, area pattern, homogeneity, artifacts, additional structures

However, Take et al. [24] have noticed that there is still a lack of possibilities to model business processes, which include the ML models. Hence, an extension of process patterns is needed. Davies et al. [7] claim that the ML use significantly increases processing speed, optimizes the processes, reduces costs, and increases purity.

III. CONCLUSION

The article presents a model of the histopathological specimen processing process, with emphasis on the part of the process concerning the processing of the digital image of the specimen. A two-stage architecture of the image processing system has been proposed, which enables automatic execution of individual tasks of the process and automation of the decision-making process regarding the image quality of the digital specimen. The proposed process model raises a number of potential possibilities. Firstly, it allows you to automate individual tasks of the process. Secondly, it reduces the burden on the pathologist by eliminating the need to deal with incorrect digital images of slides. Thirdly, it is potentially a step towards full automation of the process and the use of the Internet of Things (IoT) architecture. Authors argue that the pathomorphological diagnosis unit (PDU) procedures and processes are valuable for its control, auditing, monitoring, and digitalization and automation. The procedures for control and auditing purposes should be supported by PD process models in a

standard notation, e.g., BPMN. In this study, authors present just three processes including one in another, but the hierarchy of all PDU processes would be needed for the PDU auditing. So far, research works have not described sufficiently well the quantitative relationships between the measurable features of the PD process and the tissue effect. The global recommendations for quality control in pathomorphological diagnosis laboratories emphasize "process quality control"; however, there is no recommended model of this control, recommended control methods and "points" relevant to process quality.

The in-depth interviewing allows for the conclusions that there are also no precise indications and recommendations for dealing with problems arising during the histological process, in particular indication of the precise algorithms of actions, which depend on the type of tissue, size of the section and the correlation of the nature of artifacts (or damage) with the process disturbance. The assessment of the correctness of the process is currently based only on non-parametric evaluation - literally "by eye" of the pathomorphologist, depending on staff knowledge and experience, there are no tools for synthetic assessment of the quality of the histological process that could indicate the risk of a disorder in subsequent repetitions of the process.

Arrangement of synthetic characteristics and their definition, and establishing of the relationship between changes at the earlier stages of the process and the final

effect will enable the creation of a system of continuous quality control of the histological process in pathology departments. This study is to fill that gap, by emphasizing that the first step in PDU process audit, i.e., PDU process modeling needs to be done. In the presented models authors highlight what equipment is to be used as well as what data and algorithms would be utilized.

REFERENCES

- [1] M.M. Abdelsamea, U. Zidan, Z. Senousy, M.M. Gaber, E. Rakha, M. Ilyas, "A survey on artificial intelligence in histopathology image analysis," in *Wiley Interdisciplinary Reviews Data Mining and Knowledge Discovery* vol. 12, issue 6, pp. 1-44, 2022. <https://doi.org/10.1002/widm.1474>.
- [2] O.G. Baker, "Process surveillance: an epidemiologic challenge for all health care organizations," in *American Journal of Infection Control*, Vol. 25, Issue 2, pp. 97-101, 1997. doi: 10.1016/s0196-6553(97)90034-1. PMID: 9113284.
- [3] S. Bharati, M. R. H. Mondal, P. Podder, and V. B. S. Prasath, "Deep Learning for Medical Image Registration: A Comprehensive Review," in *International Journal of Computer Information Systems and Industrial Management Applications*, vol. 14, pp. 173-190, 2022.
- [4] *Business Process Model and Notation (BPMN) version 2.0*, Object Management Group (OMG), 2011, <https://www.omg.org/spec/BPMN/2.0/PDF/>
- [5] F. C. Dargam, E. Perz, S. Bergmann, E. Rodionova, P. Sousa, F.A. Souza, T. Matias, J. M. Ortiz, A. Esteve-Nuñez, P. Rodenas, and P.Z. Bonachela, "Operational Decision-Making on Desalination Plants: From Process Modelling and Simulation to Monitoring and Automated Control With Machine Learning," in *International Journal of Decision Support System Technology (IJDSST)*, 15(2), 2023, pp. 1-20. <http://doi.org/10.4018/IJDSST.315639>
- [6] J. De Matos, S. T. M., Ataky, A. de S. Britto, L. E. S., de Oliveira, and A.L. Koerich, "Machine learning methods for histopathological image analysis: A review," in *Electronics* (Switzerland), vol. 10, no. 5, 2021. doi: 10.3390/electronics10050562.
- [7] W. George Davies, S. Babamohammadi, Y. Yang, and S. Masoudi Soltani, "The rise of the machines: A state-of-the-art technical review on process modelling and machine learning with hydrogen production with carbon capture," in *Gas Science and Engineering*, Vol. 119, No 205104, 2023. doi: 10.1016/j.jgsce.2023.205104
- [8] M. Gholinejad, A. J. Loeve, and J. Dankelman, "Surgical process modelling strategies: which method to choose for determining workflow?" in *Minimally Invasive Therapy and Allied Technologies*, 28(2), pp. 91-104, 2019. doi: 10.1080/13645706.2019.1591457
- [9] C. Gisske, J. Liu, and K. Gand, "Applying Goal-Oriented Modelling for Machine Learning Based Rehabilitation Care," in *Studies in Health Technology and Informatics*, 294, pp. 342-346, 2022.
- [10] M.N. Gurcan, L.E. Boucheron, A. Can, A. Madabhushi, N.M. Rajpoot, and B. Yener, "Histopathological Image Analysis: A Review," in *IEEE Reviews in Biomedical Engineering*, vol. 2, pp. 147-171, 2009. doi: 10.1109/RBME.2009.2034865.
- [11] O. Jimenez-Del-Toro, S. Otalora, M. Andersson, K. Euren, M. Hedlund, M. Rousson, H. Muller, and M. Atzori, "Analysis of histopathology images: From traditional machine learning to deep learning" in *Biomedical Texture Analysis: Fundamentals, Tools and Challenges*, pp. 281-314, 2017. doi: 10.1016/B978-0-12-812133-7.00010-7.
- [12] M. Komorowski, D.C. Marshall, J.D. Saliccioli, and Y. Crutain, "Exploratory Data Analysis," in *Secondary Analysis of Electronic Health Records*. M. Komorowski, D.C. Marshall, J.D. Saliccioli, and Y. Crutain. Eds. Cham: Springer, 2016, pp. 185-204. https://doi.org/10.1007/978-3-319-43742-2_15
- [13] S. Kosaraju, J. Park, H. Lee, J.W. Yang, and M. Kang, "Deep learning-based framework for slide-based histopathological image analysis," in *Scientific Reports* 12, Nature Portfolio, No 19075, 2022. <https://doi.org/10.1038/s41598-022-23166-0>.
- [14] W. Kratsch, J. Manderscheid, M. Röglinger, and J. Seyfried, "Machine Learning in Business Process Monitoring: A Comparison of Deep Learning and Classical Approaches Used for Outcome Prediction," in *Business & Information Systems Engineering*, 63, pp. 261-276, 2021.
- [15] L. Maas, M. Geurtsen, F. Nouwt, S. Schouten, R. van der Water, S. van Dulmen, F. Dalpiaz, K. van Deemter, S. Brinkkemper, "The Care2Report System: Automated Medical Reporting as an Integrated Solution to Reduce Administrative Burden in Healthcare," in *Proceedings of the 53rd Hawaii International Conference on System Sciences*, pp. 3608-3617, 2020. Last accessed 2023/10/25 <http://hdl.handle.net/10125/64184>.
- [16] D. Marlon, M. La Rosa, J. Mendling, and H.A. Reijers, *Fundamentals of Business Process Management*, Berlin, Springer, 2018.
- [17] A. Marszałek, W. Grakowska, J. Szumiło, E. Kaznowska, Ł. Szylberg, and J. Szpor, *Podręcznik wdrożeniowy, Jednostki diagnostyki patomorfologicznej, Wskazówki i zalecenia dla jednostek diagnostyki patomorfologicznej ułatwiające wdrożenie standardów akredytacyjnych*. Kraków, Centrum Monitorowania Jakości w Ochronie Zdrowia, 2022.
- [18] D. Mouromtsev, "Semantic reference model for individualization of information processes in IoT heterogeneous environment," in *Electronics* (Switzerland), 10(20), 2523, 2021.
- [19] *Obwieszczenie Ministra Zdrowia z dnia 24 września 2021 r w sprawie standardów akredytacyjnych w zakresie udzielania świadczeń zdrowotnych oraz funkcjonowania jednostek diagnostyki patomorfologicznej*. Warszawa, Ministerstwo Zdrowia, 2021.
- [20] V. Pereira, M.P. Basilio, and C.H.T. Santos, "pyBibX--A Python Library for Bibliometric and Scientometric Analysis Powered with Artificial Intelligence Tools," arXiv preprint arXiv:2304.14516, 2023. <https://doi.org/10.48550/arXiv.2304.14516>
- [21] M. Roy, F., Wang, G. Teodoro, S. Bhattarai, M. Bhargava, T.S. Rekha, R. Aneja, and J. Kong, "Deep learning based registration of serial whole-slide histopathology images in different stains" in *Journal of Pathology Informatics* vol. 14, 2023. doi: 10.1016/j.jpi.2023.100311.
- [22] *Rozporządzenie Ministra Zdrowia z dnia 18 grudnia 2017 w sprawie standardów organizacyjnych opieki zdrowotnej w dziedzinie patomorfologii*. Dziennik Ustaw Rzeczypospolitej Polskiej poz 2435, 2017.
- [23] R. Sun, T. Lei, Q. Chen, Z. Wang, X. Du, and W. Zhao, and A.K. Nandi, "Survey of Image Edge Detection," in *Frontiers in Signal Processing*, vol. 2, Mar. 2022, doi: 10.3389/FRSIP.2022.826967.
- [24] M. Take, C. Becker, S. Alpers, and A. Oberweis, "Modeling the Integration of Machine Learning into Business Processes with BPMN," in *Proceedings of Eighth International Congress on Information and Communication Technology. ICICT 2023. Lecture Notes in Networks and Systems, vol 696*, H.S. Yang, R.S. Sherratt, N. Dey, and A. Joshi, Eds. Singapore, Springer, 2024, pp. 943-958. https://doi.org/10.1007/978-981-99-3236-8_76
- [25] J. Tang, L. Li, Y. Liu, and K.-Y. Lin, "Automatic identification of bottleneck tasks for business process management using fusion-based text clustering," in *IFAC-PapersOnLine*, 54(1), pp. 1200-1205, 2021.
- [26] X. Wang, S. Yang, J. Zhang, M. Wang, J. Zhang, W. Yang, J. Huang, and X. Han, "Transformer-based unsupervised contrastive learning for histopathological image classification," in *Medical Image Analysis*, Vol 81, 2022. doi: 10.1016/j.media.2022.102559.
- [27] W.M.P. Van der Aalst, "Process Mining: A 360 Degree Overview," in *Process Mining Handbook*, W.M.P. Van der Aalst, J. Carmona Eds. Cham, Springer, 2022, pp. 3-17.
- [28] R. Velioglu, J.P. Gopfert, A. Artelt, and B. Hammer, "Explainable Artificial Intelligence for Improved Modeling of Processes," in *23rd International Conference on Intelligent Data Engineering and Automated Learning, IDEAL 2022. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics, Vol 12756 LNCS*, H. Yin, D. Camacho, P. Tino, Eds. Cham, Springer, 2022, pp. 313-325.
- [29] Y. Xu, Z. Jia, L.B. Wang, Y. Ai, F. Zhang, M. Lai, and E.-I.-C. Chang, "Large scale tissue histopathology image classification, segmentation, and visualization via deep convolutional activation features," in *BMC Bioinformatics*, 18, 281, 2017. <https://doi.org/10.1186/s12859-017-1685-x>.

Extension-principle-based Solution Algorithm to Full LR-fuzzy Linear Programming Problems

Bogdana Stanojević
0000-0003-4524-5354

First: Mathematical Institute of the
Serbian Academy of Sciences and Arts,
Kneza Mihaila 36, 11000 Belgrade, Serbia
bgdnpop@mi.sanu.ac.rs

Second: University of Belgrade,
Faculty of Organizational Sciences,
Jove Ilića 154, 11000 Belgrade, Serbia
bogdana.stanojevic@fon.bg.ac.rs

Milan Stanojević, Nebojša Nikolić
0000-0002-7898-3401, 0000-0003-3022-5617,
University of Belgrade,
Faculty of Organizational Sciences,
Jove Ilića 154, 11000 Belgrade, Serbia
{milan.stanojevic, nebojsa.nikolic}@fon.bg.ac.rs

Abstract—In the literature one can find various methods for solving full fuzzy linear programming problems. Very few of them fully comply to the extension principle. In the current study we extend an existing solution approach based on the extension principle to derive fuzzy-set optimal results to full fuzzy linear programming problems with L-R fuzzy parameters.

Our approach is a twofold extension of a procedure from the literature: (i) it employs L-R membership functions in the optimization models that derive fuzzy-set optimal objective values; and (ii) it introduces new optimization models for deriving fuzzy-set optimal solution values in accordance to the extension principle and product operator. The employment of the product operator makes the derived fuzzy-set results more thinner, thus more appropriate to further decision making.

I. INTRODUCTION

FUZZY SETS are currently used within the soft computing field to handle uncertainties when modeling real systems. Full fuzzy programming problems are widely addressed in the recent literature. The majority of the solution approaches involve cumbersome case analyses, and provide trivial solutions as soon as the appropriate case is identified. Many times, such solutions are non-consistent, since the applied methodologies do not comply to the extension principle (EP) [1].

Particularly, procedures for solving full fuzzy linear programming problems, that work in full accordance to the extension principle, were reported in the literature. The α -cut intervals were used in all these procedures to derive the fuzzy set optimal solutions. The min operator firstly used within the extension principle to aggregate the parameters was later on replaced by the product operator in order to achieve thinner fuzzy set optimal solutions to the same problems [2].

An extension-principle-based methodology is available only for problems involving trapezoidal fuzzy parameters so far, thus the purpose of this study is to provide one for problems involving general L-R fuzzy parameters. The proposed methodology will extend the existing models for deriving the fuzzy set optimal values of the objective functions, and provide

new optimization models to obtain the fuzzy set optimal solution values of the decision variables.

Firstly, the paper contributes to the enlargement of the class of full fuzzy optimization problems that can be solved numerically. Secondly, it provides more detailed solutions, in the sense of describing numerically the optimal values of the decision variables, not only of the objective function. Moreover, the solution approach we propose can be widely used to validate other methodologies capacity to derive solutions in accordance to the extension principle. The results presented in this paper have theoretical foundations proved mathematically, and numerical illustrations.

The rest of the paper is organized as follows: Section II briefly surveys the relevant literature; Section III explains all necessary notation and provides the basic terminology; Section IV presents our solution approach. An illustrative example is offered in Section V; and the final remarks are included in Section VI.

II. LITERATURE REVIEW

Zimmermann [3] was the first to apply the fuzzy set theory to mathematical programming. A wide survey of the papers from the literature providing various models and solutions to fuzzy linear programming problems can be found in [4]. Perez-Canedo, Verdegay and Concepcion-Morales [5] reviewed the recent approaches to fuzzy linear programming based on lexicographic methods.

Full fuzzy linear programming (FF-LP) problems refer to linear-programming-shape problems having fuzzy parameters and decision variables. Majority of papers addressing FF-LP problems a priori impose the type of fuzzy sets used for both parameters and variables; apply the fuzzy arithmetic via the extension principle; and then optimize a single or multiple crisp objective function to derive the final solution. For instance Kumar, Kaur and Singh [6] imposed equality constraints on the FF-LP model, and transformed it to a crisp optimization problem using a ranking function. Later

on, Ezzati, Khorram and Enayati [7] used a lexicographic method to solve the multiple objective linear programming problem attached to the original FF-LP problem. Liu and Kao [8] introduced the solution concept to fuzzy transportation problems based on extension principle. Their approach was further adapted and improved in [9], [10].

Solutions to LP problems with trapezoidal fuzzy coefficients and using score functions were provided by Suriyapriya, Murugan, and Nayagam in [11]; and some advanced techniques to solving intuitionistic fuzzy multiple objective non-linear optimization problems were introduced by Rani, Ebrahimnejad, and Gupta in [12]. Perez-Canedo and Concepcion-Morales [13] and [14] addressed FF-LP problems with L-R fuzzy and intuitionistic fuzzy parameters and decision variables. Both their mathematical models had inequality constraints. In the second study unrestricted (free) variables were also considered. Our approach do not involve any comparison of L-R fuzzy numbers, thus avoiding any use of ranking functions and being applicable to solve any FF-LP model including any constraint types or variables. The first attempt to empirically solve FF-LP problems based on the extension principle and using the generalized "min" operator was made by Stanojevic and Nadaban [2]. They proved that by using the product operator within the extension principle more narrow fuzzy set solutions can be derived.

Summing up, the solving methods for full fuzzy optimization problems use either scalar or vector defuzzification approaches before optimization, or the extension principle without any defuzzification. Ranking functions are used to attain defuzzifications (see for instance [15] and [16]). Vector defuzzification methods use more ranking functions to transform a fuzzy optimization problem into a crisp multiple objective optimization problem. The methods that do not involve any defuzzification use α -cuts, $\alpha \in [0, 1]$, thus numerically construct the final fuzzy set solution. A summary of the literature review is provided in Table I.

III. PRELIMINARIES

The basic terminology indispensable to this study is related to L-R fuzzy membership functions (introduced by Dubois and Prade [20]) and the extension principle (introduced by Zadeh [21]).

A fuzzy set \tilde{A} of a universe U is formally described by the set of pairs $\left\{ \left(x, \mu_{\tilde{A}}(x) \right) \mid x \in U \right\}$, where $\mu_{\tilde{A}} : U \rightarrow [0, 1]$ is the membership function that applied on any element $x \in U$ provides its degree of fuzziness.

A fuzzy number is a special case of a convex, normalized fuzzy set of the universe of real numbers. An L-R fuzzy number \tilde{w}^{LR} , also referred as quadruple (w^1, w^2, w^3, w^4) , has a membership function of the following form:

$$\mu_{\tilde{A}}(x) = \begin{cases} L\left(\frac{w^2-x}{w^2-w^1}\right), & x \leq m, \\ R\left(\frac{x-w^3}{w^4-w^3}\right), & x \geq n, \\ 0, & , \end{cases}$$

where both L, R are defined on the set of real numbers such that $\mu_{\tilde{A}}(x) \in [0, 1]$. For our numerical illustrations we successively use $\max\{0, 1-x^2\}$ and $\max\{0, 1-\sqrt{x}\}$ to define both L and R functions.

A general full fuzzy optimization problem

$$\max \{f(\tilde{a}, \tilde{x}) \mid g(\tilde{c}, \tilde{x}) \leq 0\}$$

with \tilde{a} and \tilde{c} vectors of fuzzy parameters, and \tilde{x} vector of fuzzy decision variables can be solved via Zadeh's extension principle using the solution concept that computes the membership degree of the optimal solutions to the crisp optimization problem

$$\max \{f(a, x) \mid g(c, x) \leq 0\} \quad (1)$$

using the membership degree of the parameters. Let p denote the vector obtained by concatenating the vectors a and c , i.e. $p = (a|c)$. Then, the membership degree of p is defined as

$$\mu(p) = \min \left\{ \mu_{p_i}(p_i) \mid i = 1, \dots, m \right\}, \quad (2)$$

where m is the number of scalar parameters used in modeling the original optimization problem. Then, the membership degree $\mu_z(z)$, of the crisp optimal value z in the fuzzy set of the optimal values $\tilde{z} = \max \{f(\tilde{a}, \tilde{x}) \mid g(\tilde{c}, \tilde{x}) \leq 0\}$ was defined by

$$\sup \{ \mu(p) \mid p = (a|c), z = \max \{f(a, x) \mid g(c, x) \leq 0\} \} \quad (3)$$

if there exist the parameter vectors a and c such that the optimal value of Problem (1) equals to z , or $\mu_z(z) = 0$, otherwise. Similarly, for each scalar decision variable \tilde{x}_h the membership degree $\mu_{x_h}(x_h)$ is defined as

$$\sup \{ \mu(p) \mid p = (a|c), x_h = \arg_h \max \{f(a, x) \mid g(c, x) \leq 0\} \} \quad (4)$$

if there exist the parameter vectors a and c such that the h -th component of the optimal solution to Problem (1) is x_h , or $\mu_{x_h}(x_h) = 0$, otherwise.

IV. PROPOSED METHODOLOGY

The mathematical model of a FF-LP problem with L-R fuzzy parameters is given in (5).

$$\begin{aligned} \max \quad & \sum_{j=1}^n \tilde{c}_j^{LR} \tilde{x}_j, \\ \text{s.t.} \quad & \sum_{j=1}^n \tilde{a}_{ij}^{LR} \tilde{x}_j \leq \tilde{b}_i^{LR}, \quad i = \overline{1, m}, \\ & \tilde{x}_j \geq 0, \quad j = \overline{1, n}, \end{aligned} \quad (5)$$

where $\overline{1, n}$ stands for the set of natural numbers from 1 to n , i.e. $\{1, 2, \dots, n\}$. This model has non-negative decision variables and inequality constraints but it can be considered general from the point of view of our solution approach: to derive the final solutions, our approach collects the optimal results of the crisp LP problems

TABLE I
FEATURES OF OUR APPROACH AND SEVERAL APPROACHES FROM THE LITERATURE

References	Coefficients	Variables	Constraints	Optimization criterion / method
[17] Allahviranloo et al. (2008)	L-R	unrestricted	inequality	ranking function, scalar defuzzification
[6] Kumar et al. (2011)	TFN	restricted	equality	ranking function, scalar defuzzification
[7] Ezzati et al. (2015)	TFN	restricted	equality	lexicographic, vector defuzzification
[18] Kaur & Kumar (2016)	L-R	unrestricted	equality	lexicographic, vector defuzzification
[13] Perez-Canedo et al. (2019)	L-R	unrestricted	inequality	lexicographic, vector defuzzification
[19] Stanojevic & Stanojevic (2020)	TFN	unrestricted	any	EP based, "min" op., empiric z
[2] Stanojevic & Nadaban (2023)	TFN	unrestricted	any	EP based, "prod" op., empiric z
Current study	L-R	unrestricted	any	EP based, "prod" op., numeric x and z

$$\begin{aligned}
& \max \sum_{j=1}^n c_j x_j, \\
& \text{s.t.} \sum_{j=1}^n a_{ij} x_j \leq b_i, \quad i = \overline{1, m}, \\
& \quad \quad \quad x_j \geq 0, \quad \quad \quad j = \overline{1, n}.
\end{aligned} \tag{6}$$

Problem (6) is the crisp analog of (5) in the sense of having the same shape but crisp parameters; and any crisp LP problem can be converted to (6) without losing its generality.

On the other side, Problem (5) cannot represent the general form for other approaches that perform inconsistent fuzzy arithmetic operations on its parameters and variables.

Let us denote by

$$X_{A,b} = \{x \in R^m \mid A^T x \leq b^T, x \geq 0\} \tag{7}$$

the feasible set of Problem (6), and by

$$U_{A,c} = \{u \in R^m \mid A^T u \geq c^T, u \geq 0\} \tag{8}$$

the feasible set of the dual of Problem (6).

The following Theorem 4.1 presents the theoretical foundation of Algorithm 1.

Theorem 4.1: For any value α^* arbitrary fixed in the interval $[0, 1]$, the left and right endpoints of the interval representing the α^* -cut of the fuzzy set of optimal values to Model (5) are equal to the optimal values of Model (9)

$$\begin{aligned}
& \max \sum_{j=1}^n c_j x_j, \\
& \text{s.t.} \left(\prod_{i=1}^m \prod_{j=1}^n \delta_{ij} \right) \left(\prod_{i=1}^m \beta_i \right) \left(\prod_{j=1}^n \gamma_j \right) = \alpha^*, \\
& \quad a_{ij}^3 L_{a_{ij}^{LR}}^{-1}(\delta_{ij}) - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 R_{a_{ij}^{LR}}^{-1}(\delta_{ij}) + a_{ij}^2, \\
& \quad b_i^3 L_{b_i^{LR}}^{-1}(\beta_i) - b_i^1 \leq b_i \leq b_i^4 R_{b_i^{LR}}^{-1}(\beta_i) + b_i^2, \\
& \quad c_j^3 L_{c_j^{LR}}^{-1}(\gamma_j) - c_j^1 \leq c_j \leq c_j^4 R_{c_j^{LR}}^{-1}(\gamma_j) + c_j^2, \\
& \quad \delta_{ij}, \beta_i, \gamma_j \in [0, 1], i = \overline{1, m}, j = \overline{1, n}, \\
& \quad x \in X_{A,b};
\end{aligned} \tag{9}$$

and Model (10)

$$\begin{aligned}
& \min \sum_{i=1}^m b_i u_i, \\
& \text{s.t.} \left(\prod_{i=1}^m \prod_{j=1}^n \delta_{ij} \right) \left(\prod_{i=1}^m \beta_i \right) \left(\prod_{j=1}^n \gamma_j \right) = \alpha^*, \\
& \quad a_{ij}^3 L_{a_{ij}^{LR}}^{-1}(\delta_{ij}) - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 R_{a_{ij}^{LR}}^{-1}(\delta_{ij}) + a_{ij}^2, \\
& \quad b_i^3 L_{b_i^{LR}}^{-1}(\beta_i) - b_i^1 \leq b_i \leq b_i^4 R_{b_i^{LR}}^{-1}(\beta_i) + b_i^2, \\
& \quad c_j^3 L_{c_j^{LR}}^{-1}(\gamma_j) - c_j^1 \leq c_j \leq c_j^4 R_{c_j^{LR}}^{-1}(\gamma_j) + c_j^2, \\
& \quad \delta_{ij}, \beta_i, \gamma_j \in [0, 1], i = \overline{1, m}, j = \overline{1, n}, \\
& \quad u \in U_{A,c},
\end{aligned} \tag{10}$$

respectively.

Within these models, the objective functions are optimized over the variables $a_{ij}, b_i, c_j, \delta_{ij}, \beta_i, \gamma_j, i = \overline{1, m}, j = \overline{1, n}$. To derive the left (right) endpoints of the α^* -cut of the fuzzy-set optimal solutions to Problem (5), for any fixed index $j^* \in \{1, 2, \dots, n\}$, and $\alpha^* \in [0, 1]$, we introduce Model (11)

$$\begin{aligned}
& \min(\max) x_{j^*}, \\
& \text{s.t.} \\
& \quad z_{\alpha^*}^{\min} \leq \sum_{j=1}^n c_j x_j \leq z_{\alpha^*}^{\max}, \\
& \quad \left(\prod_{i=1}^m \prod_{j=1}^n \delta_{ij} \right) \left(\prod_{i=1}^m \beta_i \right) \left(\prod_{j=1}^n \gamma_j \right) = \alpha^*, \\
& \quad a_{ij}^3 L_{a_{ij}^{LR}}^{-1}(\delta_{ij}) - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 R_{a_{ij}^{LR}}^{-1}(\delta_{ij}) + a_{ij}^2, \\
& \quad b_i^3 L_{b_i^{LR}}^{-1}(\beta_i) - b_i^1 \leq b_i \leq b_i^4 R_{b_i^{LR}}^{-1}(\beta_i) + b_i^2, \\
& \quad c_j^3 L_{c_j^{LR}}^{-1}(\gamma_j) - c_j^1 \leq c_j \leq c_j^4 R_{c_j^{LR}}^{-1}(\gamma_j) + c_j^2, \\
& \quad \delta_{ij}, \beta_i, \gamma_j \in [0, 1], i = \overline{1, m}, j = \overline{1, n}, \\
& \quad x \in X_{A,b},
\end{aligned} \tag{11}$$

where $z_{\alpha^*}^{\min}$ and $z_{\alpha^*}^{\max}$ are the optimal values obtained by solving Models (9) and (10) for the given value $\alpha^* \in [0, 1]$.

Algorithm 1 formally describes our solution approach. Generally, the predefined values $\alpha_1^*, \alpha_2^*, \dots, \alpha_q^*$ are chosen to be equidistant in the interval $[0, 1]$, $\alpha_1^* = 0, \alpha_q^* = 1$.

Algorithm 1 Deriving the fuzzy set solutions

Require: the predefined α -levels $\alpha_1^*, \alpha_2^*, \dots, \alpha_q^*$, and the L-R fuzzy-parameter matrices $\tilde{A}^{LR}, \tilde{b}^{LR}, \tilde{c}^{LR}$.

Ensure: $z^{\min}, z^{\max}, x^{\min}, x^{\max}$.

for $k = 1, q$ **do**

2: Set $\alpha^* = \alpha_k^*$, and derive $z_{\alpha_k^*}^{\min}$ and $z_{\alpha_k^*}^{\max}$ by solving (9) and (10), respectively.

for $h = 1, n$ **do**

4: Derive $x_{\alpha_k^*}^{\min}(h^*)$ and $x_{\alpha_k^*}^{\max}(h^*)$ as minimal and maximal values of the objective function in (11).

end for

6: Construct the vectors $x_{\alpha_k^*}^{\min} = \left(x_{\alpha_k^*}^{\min}(h^*) \right)_{h^*=\overline{1,n}}$ and

$$x_{\alpha_k^*}^{\max} = \left(x_{\alpha_k^*}^{\max}(h^*) \right)_{h^*=\overline{1,n}}.$$

end for

8: Construct $z^{\min} = \left(z_{\alpha_k^*}^{\min} \right)^{k=\overline{1,q}}$, $z^{\max} = \left(z_{\alpha_k^*}^{\max} \right)^{k=\overline{1,q}}$,

$$x^{\min} = \left(x_{\alpha_k^*}^{\min} \right)^{k=\overline{1,q}}, x^{\max} = \left(x_{\alpha_k^*}^{\max} \right)^{k=\overline{1,q}}.$$

A. Particularities due to specific shapes of L-R functions

In what follows we discuss the particularities of Models (9), (10) and (11) in certain cases, with respect to the L-R functions that are commonly used. All particular models are polynomial, or equivalent to polynomial models.

1) $L(x) = R(x) = \max\{0, 1 - x\}$: This particular case corresponds to trapezoidal fuzzy numbers. The fuzzy set of the optimal values of the objective functions were empirically described in [19] using min operator; and numerically derived in [2] using product operator. The new introduced Model (11) adapted to the current case becomes

$$\min(\max) \quad x_{j^*},$$

s.t.

$$z_{\alpha^*}^{\min} \leq \sum_{j=1}^n c_j x_j \leq z_{\alpha^*}^{\max},$$

$$\left(\prod_{i=1}^m \prod_{j=1}^n \delta_{ij} \right) \left(\prod_{i=1}^m \beta_i \right) \left(\prod_{j=1}^n \gamma_j \right) = \alpha^*,$$

$$a_{ij}^3 (1 - \delta_{ij}) - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 (1 - \delta_{ij}) + a_{ij}^2,$$

$$b_i^3 (1 - \beta_i) - b_i^1 \leq b_i \leq b_i^4 (1 - \beta_i) + b_i^2,$$

$$c_j^3 (1 - \gamma_j) - c_j^1 \leq c_j \leq c_j^4 (1 - \gamma_j) + c_j^2,$$

$$\delta_{ij}, \beta_i, \gamma_j \in [0, 1], i = \overline{1, m}, j = \overline{1, n},$$

$$x \in X_{A,b}.$$

(12)

2) $L(x) = R(x) = \max\{1 - \sqrt{x}\}$: In this case, the inverse functions of L and R are $L^{-1}(y) = R^{-1}(y) = (1 - y)^2$, and the constraints on variables $a_{ij}, b_i, c_j, i = \overline{1, m}, j = \overline{1, n}$ become quadratic, i.e.

$$a_{ij}^3 (1 - \delta_{ij})^2 - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 (1 - \delta_{ij})^2 + a_{ij}^2,$$

$$b_i^3 (1 - \beta_i)^2 - b_i^1 \leq b_i \leq b_i^4 (1 - \beta_i)^2 + b_i^2, \quad (13)$$

$$c_j^3 (1 - \gamma_j)^2 - c_j^1 \leq c_j \leq c_j^4 (1 - \gamma_j)^2 + c_j^2,$$

3) $L(x) = R(x) = \max\{1 - x^2\}$: In this case, the inverse functions of L and R are $L^{-1}(y) = R^{-1}(y) = \sqrt{1 - y}$,

TABLE II
FUZZY PARAMETERS OF PROBLEM (15)

Objective function and RHS	Constraints matrix
$c_1 = (6, 8, 8, 10)$	$\tilde{a}_{11} = (4.5, 5, 5, 5.5)$
$\tilde{c}_2 = (10, 12, 12, 14)$	$\tilde{a}_{21} = (5.75, 6, 6, 6.25)$
$c_3 = (0.75, 1, 1, 1.25)$	$a_{31} = (0.5, 1, 1, 1.25)$
$\tilde{b}_1 = (105, 150, 155, 207)$	$\tilde{a}_{12} = (4.5, 5, 5, 5.5)$
$\tilde{b}_2 = (102, 120, 125, 147)$	$\tilde{a}_{22} = (1.75, 2, 2, 2.25)$
$\tilde{b}_3 = (58, 100, 110, 148)$	$\tilde{a}_{32} = (3.75, 4, 4, 4.5)$

and the constraints on variables $a_{ij}, b_i, c_j, i = \overline{1, m}, j = \overline{1, n}$ become

$$\begin{aligned} a_{ij}^3 \sqrt{1 - \delta_{ij}} - a_{ij}^1 &\leq a_{ij} \leq a_{ij}^4 \sqrt{1 - \delta_{ij}} + a_{ij}^2, \\ b_i^3 \sqrt{1 - \beta_i} - b_i^1 &\leq b_i \leq b_i^4 \sqrt{1 - \beta_i} + b_i^2, \\ c_j^3 \sqrt{1 - \gamma_j} - c_j^1 &\leq c_j \leq c_j^4 \sqrt{1 - \gamma_j} + c_j^2. \end{aligned} \quad (14)$$

With the help of the transformations

$$\begin{aligned} \eta_{ij} &= \sqrt{1 - \delta_{ij}}, \theta_i = \sqrt{1 - \beta_i}, \\ \vartheta_j &= \sqrt{1 - \gamma_j}, i = \overline{1, m}, j = \overline{1, n}, \end{aligned}$$

constraints (14) become linear, and the degree of the polynomial in the second constraint of Models (9), (10) and (11) is doubled. As a consequence, the final equivalent constraint system

$$z_{\alpha^*}^{\min} \leq \sum_{j=1}^n c_j x_j \leq z_{\alpha^*}^{\max},$$

$$\left(\prod_{i=1}^m \prod_{j=1}^n (1 - \eta_{ij}^2) \right) \left(\prod_{i=1}^m (1 - \theta_i^2) \right) \left(\prod_{j=1}^n (1 - \vartheta_j^2) \right) = \alpha^*,$$

$$a_{ij}^3 \eta_{ij} - a_{ij}^1 \leq a_{ij} \leq a_{ij}^4 \eta_{ij} + a_{ij}^2,$$

$$b_i^3 \theta_i - b_i^1 \leq b_i \leq b_i^4 \theta_i + b_i^2,$$

$$c_j^3 \vartheta_j - c_j^1 \leq c_j \leq c_j^4 \vartheta_j + c_j^2,$$

$$\eta_{ij}, \theta_i, \vartheta_j \in [0, 1], i = \overline{1, m}, j = \overline{1, n},$$

$$x \in X_{A,b},$$

is polynomial.

V. ILLUSTRATIVE EXAMPLE

Perez-Canedo et al. [13] used the following fuzzy optimization problem (15) to illustrate their approach.

$$\begin{aligned} \max \quad & \tilde{c}_1 \tilde{x}_1 + \tilde{c}_1 \tilde{x}_2 + \tilde{c}_1 \tilde{x}_3 \\ \text{s.t.} \quad & \tilde{a}_{11} \tilde{x}_1 + \tilde{a}_{12} \tilde{x}_2 + \tilde{x}_3 = \tilde{b}_1, \\ & \tilde{a}_{21} \tilde{x}_1 + \tilde{a}_{22} \tilde{x}_2 \leq \tilde{b}_2, \\ & \tilde{a}_{31} \tilde{x}_1 + \tilde{a}_{32} \tilde{x}_2 \leq \tilde{b}_3, \\ & \tilde{x}_1, \tilde{x}_2 \geq 0, \\ & \tilde{x}_3 \text{ free variable.} \end{aligned} \quad (15)$$

The fuzzy parameters of Problem (15) are given in Table II. Problem (15) has both equality and inequality constraints; and both bounded and unbounded variables. These characteristics make it relevant to describe the generality of the solution approaches.

Perez-Canedo et al. solved four variants of this problem using: (i) $L(x) = R(x) = 1 - x$; (ii) $L(x) = R(x) = 1 - x^2$;




Fig. 1. Fuzzy set optimal values of the objective function

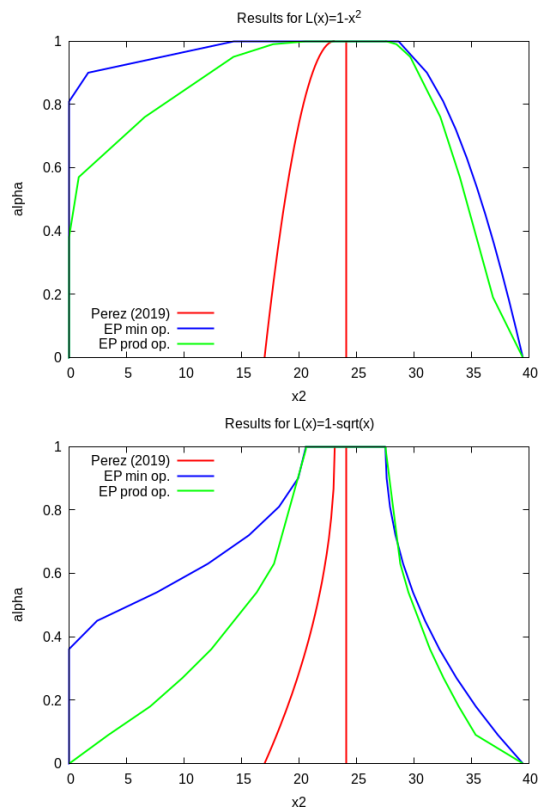


Fig. 3. Fuzzy set optimal values of variable x_2




Fig. 2. Fuzzy set optimal values of variable x_1




Fig. 4. Fuzzy set optimal values of variable x_3

(iii) $L(x) = R(x) = 1 - \sqrt{x}$; (iv) $L(x) = 1 - \sqrt{x}$ and $R(x) = 1 - x^2$.

The graphic representations obtained by applying our approach in comparison with the results obtained by Perez-Canedo et al. [13] are shown in Figures 1, 2 and 3. Figure 1 presents the fuzzy sets of the optimal objective values for both cases (ii) and (iii). Figures 2 and 3 present the fuzzy sets of the optimal solutions for the cases (ii) and (iii), respectively.

There are several facts that are well illustrated by this results: (i) the support of the fuzzy sets results (optimal objective and solution values) are not influenced by the operator “min” or “prod” used to aggregate the fuzzy parameters within extension principle, since whenever one parameter has membership degree 0, both “min” and “prod” operators provide the same membership value 0 after aggregation (ii) the shapes of the fuzzy set results (optimal objective and solution values) are much thinner when “prod” operator is used within extension-principle-based aggregation compared to “min” operator; (iii) the results obtained by Perez-Canedo et al. [13] are L-R fuzzy numbers and ours are not. Perez-Canedo et al. derived the support of the fuzzy sets results, and then applied the a priori imposed rule for obtaining the membership functions, while we derived the left and right endpoints of each desired α -cut.

VI. CONCLUSION AND FURTHER RESEARCHES

By this study we aimed to provide a solution algorithm to full fuzzy linear programming problems with L-R fuzzy descriptions to uncertain parameters. Comparing to other studies from the literature one of the advantages of our methodology is that it fully complies to the extension principle. In addition it uses the product operator to aggregate the L-R fuzzy quantities, thus deriving more narrow fuzzy set solutions to the original problem. Our approach introduced new optimization models for deriving fuzzy-set optimal solution values, and derives results that comply to the extension principle. The proposed methodology is illustrated on a numerical example recalled from the literature. The class of fuzzy optimization problems that can be solved using similar principles can be further extended, e.g. nonlinear optimization problems in fuzzy environment can be addressed. Based on the same solution concept, an empiric variant of the new introduced approach - that simulates the extension principle by choosing randomly values of the parameters within their corresponding fuzzy sets - can be employed to estimate the fuzzy set solutions in accordance to the extension principle.

REFERENCES

- [1] B. Stanojevic, M. Stanojevic, and S. Nadaban, “Reinstatement of the extension principle in approaching mathematical programming with fuzzy numbers,” *Mathematics*, vol. 9, no. 11, 2021. doi: 10.3390/math9111272
- [2] B. Stanojević and S. Nădăban, “Empiric solutions to full fuzzy linear programming problems using the generalized “min” operator,” *Mathematics*, vol. 11, no. 23, p. 4864, 2023. doi: 10.3390/math11234864
- [3] H.-J. Zimmermann, “Applications of fuzzy set theory to mathematical programming,” *Information Sciences*, vol. 36, no. 1, pp. 29 – 58, 1985. doi: 10.1016/0020-0255(85)90025-8
- [4] R. Ghanbari, K. Ghorbani-Moghadam, and B. De Baets, “Fuzzy linear programming problems: models and solutions,” *Soft Computing*, pp. 1433–1479, 2019. doi: 10.1007/s00500-019-04519-w
- [5] B. Perez-Canedo, J. L. Verdegay, E. R. Concepcion-Morales, and A. Rosete, “Lexicographic methods for fuzzy linear programming,” *Mathematics*, vol. 8, no. 9, 2020. doi: 10.3390/math8091540
- [6] A. Kumar, J. Kaur, and P. Singh, “A new method for solving fully fuzzy linear programming problems,” *Applied Mathematical Modelling*, vol. 35, no. 2, pp. 817 – 823, 2011. doi: 10.1016/j.apm.2010.07.037
- [7] R. Ezzati, V. Khorram, and R. Enayati, “A new algorithm to solve fully fuzzy linear programming problems using the molp problem,” *Applied Mathematical Modelling*, vol. 39, no. 12, pp. 3183 – 3193, 2015.
- [8] S.-T. Liu and C. Kao, “Solving fuzzy transportation problems based on extension principle,” *European Journal of Operational Research*, vol. 153, no. 3, pp. 661 – 674, 2004. doi: 10.1016/S0377-2217(02)00731-2
- [9] G. Singh and A. Singh, “Extension of particle swarm optimization algorithm for solving transportation problem in fuzzy environment,” *Applied Soft Computing*, vol. 110, p. 107619, 2021. doi: 10.1016/j.asoc.2021.107619
- [10] M. Bisht, I. Beg, and R. Dangwal, “Optimal solution of pentagonal fuzzy transportation problem using a new ranking technique,” *Yugoslav Journal of Operations Research*, vol. 33, no. 4, pp. 509–529, 2023. doi: 10.2298/YJOR221120002B
- [11] K. Suriyapriya, J. Murugan, and V. L. G. Nayagam, “Solution of linear programming problem with trapezoidal fuzzy coefficients using score functions,” *International Journal of Mathematics in Operational Research*, vol. 22, no. 1, pp. 41–73, 2022. doi: 10.1504/IJ-MOR.2022.123126
- [12] D. Rani, A. Ebrahimnejad, and G. Gupta, “Generalized techniques for solving intuitionistic fuzzy multi-objective non-linear optimization problems,” *Expert Systems with Applications*, vol. 202, p. 117264, 2022. doi: 10.1016/j.eswa.2022.117264
- [13] B. Perez-Canedo and E. Concepcion-Morales, “A method to find the unique optimal fuzzy value of fully fuzzy linear programming problems with inequality constraints having unrestricted l-r fuzzy parameters and decision variables,” *Expert Systems with Applications*, vol. 123, pp. 256 – 269, 2019. doi: 10.1016/j.eswa.2019.01.041
- [14] B. Pérez-Cañedo and E. Concepción-Morales, “On l-r-type fully intuitionistic fuzzy linear programming with inequality constraints: Solutions with unique optimal values,” *Expert Systems with Applications*, vol. 128, pp. 246 – 255, 2019. doi: 10.1016/j.eswa.2019.03.035
- [15] D. Ewald, W. Dobrosielski, J. M. Czerniak, and H. Zarzycki, “Comparative study: Defuzzification functions and their effect on the performance of the onbee optimization algorithm,” in *Communication Papers of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 37. PTI, 2023. doi: 10.15439/2023F5828 p. 91–96.
- [16] D. Kuchta, J. Grobelny, R. Michalski, and J. Schneider, “Perception of vector and triangle representations of fuzzy number most possible value changes,” in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35. IEEE, 2023. doi: 10.15439/2023F9557 p. 269–274.
- [17] T. Allahviranloo, F. H. Lotfi, M. K. Kiasary, N. Kiani, and L. Alizadeh, “Solving fully fuzzy linear programming problem by the ranking function,” *Applied mathematical sciences*, vol. 2, no. 1, pp. 19–32, 2008.
- [18] J. Kaur and A. Kumar, “Unique fuzzy optimal value of fully fuzzy linear programming problems with equality constraints having l-r flat fuzzy numbers,” in *An Introduction to Fuzzy Linear Programming Problems: Theory, Methods and Applications*. Cham: Springer International Publishing, 2016. doi: 10.1007/978-3-319-31274-3_6 pp. 109–118.
- [19] B. Stanojević and M. Stanojević, “Empirical versus analytical solutions to full fuzzy linear programming,” in *Intelligent Methods for Computing, Communications and Control. ICC2020*, ser. Advances in Intelligent Systems and Computing, I. Dzitac, S. Dzitac, F. Filip, J. Kacprzyk, M. Manolescu, and H. Oros, Eds., vol. 1243. Springer, Cham, 2020. doi: 10.1007/978-3-030-53651-0_19
- [20] D. Dubois and H. Prade, “Operations on fuzzy numbers,” *International Journal of Systems Science*, vol. 9, no. 6, pp. 613–626, 1978. doi: 10.1080/00207727808941724
- [21] L. Zadeh, “The concept of a linguistic variable and its application to approximate reasoning i,” *Information Sciences*, vol. 8, no. 3, pp. 199 – 249, 1975. doi: 10.1016/0020-0255(75)90036-5

Strategy Registry: an optimized implementation of the Strategy design pattern in solidity for the Ethereum Blockchain

Hamza Tamenaoul*, Mahmoud El Hamlaoui[†] and Mahmoud Nassar[‡]
 RABAT IT CENTER
 ENSIAS, Mohammed V University in Rabat
 Rabat, Morocco

Email: *hamza_tamenaoul@um5.ac.ma, [†]mahmoud.elhamlaoui@um5.ac.ma, [‡]mahmoud.nassar@um5.ac.ma

Abstract—The strategy design pattern is an essential behavioral design pattern. When developing software it is relied upon for the development of modular components. Given the constraints inherent to the inner workings of the Blockchain and features offered by Smart Contract languages, some of the design patterns commonly used in software systems development cannot be naively implemented. This paper explores a new pattern implementation of the mentioned pattern in a contract oriented language aimed at Smart-Contracts development, specifically the solidity language. The pattern’s implementations discussed in this paper are intended to be deployed and run on the Ethereum Blockchain.

Index Terms—Blockchain, Ethereum, Solidity, Smart Contract, Design Pattern, Strategy pattern, Gas, Optimization

I. INTRODUCTION

THE STRATEGY design pattern [1] tries to solve a very specific behavioral problem. Component are developed to run a fixed set of algorithms. Given the static nature of compiled code, those algorithms could not be defined at a different moment than during compile time. However, in multiple cases, the algorithms to be executed are not necessarily known at that moment, the information would be rather uncovered only during runtime. Therefore a design pattern had to be defined to tackle this specific issue, which is the *Strategy Design Pattern*.

The strategy design pattern defines two elements: the context and the strategy [2]. The context is the component, containing and controlling the algorithms to be executed as well as handling and managing the set of attributes or parameters that impacts the execution flow of said algorithms. From a client perspective, the context is the entity that is called at compile time to run the algorithm.

Decoupling the context from the logic gives us the opportunity to have a modular logic, while enabling us to implement multiple algorithms with a simpler implementation and easier maintainability. This design is essential in a lot of contexts when developing software libraries or software systems. This loosely coupled structure makes it possible to create a context object instance without worrying whether or not the algorithm to be used when executing the logic is known at that exact moment. Moreover, it gives us the opportunity to switch strategies during runtime easily and in

a transparent manner. The link between the strategy and the context is created when the algorithm to be used is known, sealing the relationship between the two objects and binding the strategy object lifecycle to the context object. While the strategy object can be discarded while the context is still used, for the context, discarding a strategy would not have any impact on its lifecycle. In fact, this pattern makes the lifecycle of the strategy object completely dependent on the one of the context. This relationship is illustrated in the Fig. 1




Fig. 1. UML class diagram of the implementation of the strategy design pattern

A. Pattern client implementation

Following the blueprint given by this implementation, the client code logic is required to create a new instance of a strategy every time a new context is instantiated. The following algorithm illustrate this logic.

```
Context context = new Context()
```

The context object is creating without any reference to the strategy object and therefore the algorithm to be used when running the logic.

```
Strategy strategy =
    new StrategyImplementation()
context.addStrategy(strategy)
```

When the strategy is known during runtime, the following algorithm would create the strategy object instance and bind it

to the context. From this moment the context instance has the information on what algorithm should be run forward. This part of the algorithm can be easily executed multiple time during the lifetime of the program with different strategies to accommodate any change of the logic during runtime.

```
context.doLogic()
```

Finally the logic can be called directly from the context, regardless of what logic resides underneath. This transparency when executing the logic, that makes the strategy design pattern an essential design pattern when developing any system.

In the context of non distributed application - *referred to as App* -, this client implementation would lead to the creation of new strategy instance whenever the second part of the algorithm is executed. The footprint of this approach lays mainly on the memory, requiring it to handle the whole lifecycle of the strategy object for every new instance, which is not generally an issue in modern languages.

In such standard applications the memory trade off is worth it, limited and minimal. More generally the memory footprint is negligible in front of the maintainability and the modularity gains, in distributed applications - *commonly referred to as dApps* -, we cannot neglect the memory footprint of this client implementation for cost efficiency.

II. COST ANALYSIS OF THE STRATEGY PATTERN IMPLEMENTATION IN A DISTRIBUTED APPLICATION

A. The optimization requirements of Blockchain specific code

Distributed applications rely on the Blockchain technology to be deployed and run [4]. They use Smart-contracts to define their business logic, a distributed logic. By design, the Blockchain is a technology that distributes the execution of the logic on all of its users, referred to as nodes. This structure makes also the cost of executing some code on the infrastructure cost on all nodes, leading to a higher resource cost usage cost for the user of some dApp.

This cost is very high and cannot be negligible and therefore kept to a minimum. When running Apps on a non Blockchain environment, the resource cost is the cost of the computing time of the machine that executed the code. On a Blockchain setting, that cost is multiplied by the number of nodes on the Blockchain, since they all need to execute the code for a dApp to work. Those nodes maintain and run the infrastructure of the Blockchain, therefore the Blockchain put a price on every action on the Blockchain to compensate the nodes for their work.

To calculate this compensation, the concept of **Gas** [4] was introduced to evaluate the price equivalent of every action on its infrastructure, the inherent cost of handling a transaction or running an algorithm on the Blockchain network. Gas is a notion that is specific to the Ethereum Virtual Machine (EVM) and similar Blockchains, it represents the computational effort needed to execute a transaction on the Blockchain. When multiplied by the price of a unit of gas we can calculate the price of said transaction. Therefore gas can be used as

a measure of the cost of a transaction. To be able to run code on the Blockchain, the caller is required to provide sufficient Gas for its request, or it will not be processed by the network. This requirement forces dApps designers to build and design specifically optimized code for Blockchain applications, code that would require the least amount of memory and processing resources possible to make the cost of using the dApp minimal and acceptable for users.

B. The limitations of the standard strategy pattern implementation

The standard approach to implement the strategy design pattern falls short in memory optimization aspects. In non Blockchain based languages, instantiating new objects is equivalent to allocating a new memory space on the machine running the program. In the context of the Blockchain, creating a new instance of an object, means the creation and deployment of a new Smart-contract, which means allocating resources on all nodes of the infrastructure which is one of the heaviest operations in terms of memory allocation and instructions to execute on the Virtual Machine. This makes the operation require a lot of gas which in turn translates to an extremely high cost on the burden of users. Therefore this approach is not optimal.

III. NAME REGISTRY DESIGN PATTERN BASED OPTIMIZATION

A. Name Registry pattern

On the Ethereum Blockchain, a Smart-contract needs to be deployed first before being used. Once a Smart-contract is deployed, the only way interact with the Smart-contract is by calling through its address on the Blockchain. However loosing the address of a Smart-contract means more or less loosing the ability interact with it. To solve this issue and others related to it the **Name Registry pattern** [5] was introduced.

The Name Registry pattern as its name suggest centers on a registry as the main piece, an address registry. It stores the different addresses of the Smart-Contracts created. It solves the issue of dealing with storing the physical address of Smart-contracts by taking care of this part, and providing keys to retrieve easily the addresses that would be required. The registry pattern can be simplified as a *mapping* from and already defined key - that can be mutable - to a Smart-contract address. While adding addresses inside of the registry has a cost, it is much more cheaper than creating a new Smart-Contract every time on is needed, specially when those Contracts are stateless, but depending on the implementation, the owner of the dApp can implement it a way that the burden of the cost would not fall on the user but on the owner of the Smart-Contract.

A basic key implemented as a string identifier, as shown in Fig. 2 could be used for by mapping in the Name Registry pattern. The key could be much sophisticated depending on the requirements of the project. The most common addition of to the key is the Smart-contract version, to handle versioning which is not native to the Ethereum Blockchain. This part of




Fig. 2. UML diagram of the Name Registry pattern

the pattern will not be discussed in this paper and falls out of the scope of the paper.

B. Contract independent strategy pattern implementation

To solve the overhead issue of having to deploy a new Smart-contract every time a new one defining a strategy is deployed, both the Strategy pattern could be merged with the Name Registry pattern to provide the optimization needed for the constraints of the Blockchain.

This paper will refer to the resulting pattern as the **Strategy Registry**. The Strategy Registry optimizes the Strategy pattern by removing the direct link between the context and the strategy. Instead, the context is linked to a Name Registry that contains the different strategies deployed. Strategies are registered into the registry making them context independent and unbinding their lifecycles from the context's. While the registry would contain the strategies that have been deployed by the system's development team, the nature of Blockchain development can make the addition of new strategies more independent.

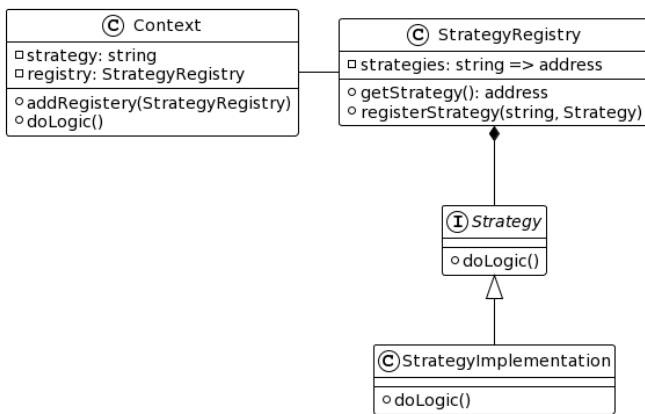


Fig. 3. UML class diagram of the Strategy Registry pattern

This optimization operates on multiple levels:

1) *Gas optimization:* With the implementation of the strategy design pattern, the number of the strategies that exists in the execution context can reach the total number of strategies defined in the system. Therefore the total number of Strategies that would coexist on the Blockchain is theoretically unlimited,

which is an important optimization issue. This would entail that for every execution an enormous amount of gas would be spent every time to repeat the same operation, with transactions containing duplicated amount of information that needs to be stored.

Using the Registry strategy pattern caps the number of strategies that could be deployed to the number of strategy implementations that were developed. Having a fixed number of contract to deploy simplifies the cost calculation and reduces enormously the overhead required during every execution. In fine, this would lead to a reduces and acceptable gas price for the pattern.

2) *Space optimization:* One of the main constraints of the Blockchain is the limited size of the available space. Distributed application designer seek to build space efficient Smart-contract, like during the first days of computing. Using the strategy design pattern breaks this rule. As mentioned previously, deploying strategies with the context only contributes to enlarging the memory size required by the context. If we take into account that the number of strategies is non capped, this means that dApp relying on such architecture could take an unlimited amount of space on the Blockchain, which is far from being ideal, nor acceptable in terms of design or cost.

The Strategy Registry solves perfectly this issues. Capping the number of Smart-contracts that could be possibly deployed means a finite amount of space would be needed for the strategies regardless of the number of executions. Moreover decoupling the Strategy from the context, leads to having lighter Smart-contracts defining the context to deploy, optimizing even further the usage of space.

3) *Versioning of strategies:* One of the main benefits of using Strategies is the ability to update the logic of some program without impacting the code of the context. This ability can be seen as lost when talking about deploying the strategies beforehand. While at fist glance this might seem true, in reality using the registry enables this feature in a more straight forward manner.

One of the main benefits of the Name Registry pattern in the Ethereum Blockchain is its ability to introduce Smart-contract versioning in a simple and easy way. This capability could be also used to version strategies or even decommission or disable them. In the naive implementation, changing or improving a strategy would have required a double cost. The first and most expensive one, is the redeployment cost of new Smart-contracts representing the strategies as many times as the old versions were deployed. The second cost is related to updating the strategies in the different Smart-contract already in production.

Using the Strategy Registry, the cost would be reduced to one of deploying the new Strategy Smart-contract and registering it in the registry alone. No other data would be changed in any other Smart-contract; the context would still use the same key - that do not have to be changed - to call the strategy, the registry makes sure that the correct Smart-contract will be called, since it holds the information about the statuses of the different Strategies.

IV. CONCLUSION

The current Blockchain technology comes with new set of constraints, constraints that makes following many of the software principles already used and theorized hard to follow. One of important principles are design patterns such the *Strategy Design Pattern*. The *Strategy Registry Design Pattern* makes such transition. It solves the problem *Strategy Design Pattern* intents solving, while providing a solution that fits the challenges faced during Ethereum Blockchain development.

While using the Strategy Registry design pattern has obvious advantages in terms of production costs for both the owner of the dApp and its user, the performance of such pattern on the Blockchain should be assessed depending on the usage. One fits all solutions rarely exist in the software world, and

even more so when working with Ethereum Smart-Contracts.

REFERENCES

- [1] E. Freeman and E. Robson, Head first design patterns: a brain-friendly guide, Second release. in A brain-friendly guide. Beijing Boston Farnham Sebastopol Tokyo: O'Reilly, 2014.
- [2] Refactoring Guru, "Strategy," Refactoring.guru, 2014. <https://refactoring.guru/design-patterns/strategy>
- [3] "Solidity — Solidity 0.8.26 documentation," docs.soliditylang.org. <https://docs.soliditylang.org/en/v0.8.26/> (accessed May 28, 2024).
- [4] B. Vitalik, "Ethereum white paper: A next-generation smart contract and decentralized application platform." 2014. [Online]. Available: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf
- [5] "Contract Registry," Blockchain Patterns. <https://research.csiro.au/blockchainpatterns/general-patterns/contract-structural-patterns/contract-registry/> (accessed May 28, 2024).

Successfully Improving the User Experience of an Artificial Intelligence System

Alexander Zender, Bernhard G. Humm, Anna Holzheuser
0000-0002-6956-9049, 0000-0001-7805-1981, 0009-0006-9376-3382
Darmstadt University of Applied Sciences
Schöfferstr. 3, 64295 Darmstadt, Germany
Email: alexander.zender@h-da.de, bernhard.humm@h-da.de, anna.holzheuser@gmail.com

Abstract—An important aspect of Artificial Intelligence (AI) Systems is their User Experience (UX), which can impact the user’s trust in the AI system. However, UX has not yet been in the focus of AI research. In previous research, we have evaluated the UX of the Meta AutoML platform OMA-ML, uncovering weak points and proposing several recommendations for ensuring a positive UX in AI systems. In this paper we show that implementing those recommendations leads to measurable UX improvements. We present the UX-improving features implemented in a new release of OMA-ML and the results from a second UX evaluation. The UX of OMA-ML could successfully be improved in four interactive principles (suitability for the user’s tasks, self-descriptiveness, user engagement and learnability). We argue that an iterative approach to UX potentially leads to more human-centered AI.

I. INTRODUCTION

ARTIFICIAL INTELLIGENCE (AI) SYSTEMS are programs or machines that can mimic human cognitive behaviour [1]. AI systems, in particular ones using Machine Learning (ML), are present in our everyday use, e.g. facial recognition in smartphones [2] or translation tools using natural language processing (NLP) [3]. An important aspect of AI systems is their User Experience (UX). UX in the context of AI systems assesses a user’s overall experience with the AI system [4]. Understanding the user’s needs and behaviours is necessary, as a bad UX may contribute to an AI system’s failure [5]. While the UX is an important aspect of an AI system, there is limited discussion about it in the AI community.

We discussed this in previous research [6], and aimed to raise awareness by using the case study for the AI platform OMA-ML¹ (Ontology-based Meta AutoML) [7][8]. Based on the case study, 104 UX issues were found, categorised and resolved. Additionally, we proposed 12 measures and 4 recommendations to ensure a positive UX for AI systems. Based on the same Methodology [6], the interaction of 29 participants using the updated version of OMA-ML was evaluated.

In this paper we show that implementing those recommendations leads to measurable UX improvements. We present the results from the new case study and the UX improvements implemented in a new release of OMA-ML. Based on the ISO 9241-110² interaction principles, OMA-ML could improve in

the interaction principles: suitability for the user’s tasks, self-descriptiveness and user engagement compared to the previous evaluations weak points and even exceed in the learnability interaction principle beyond the target state. However, it has yet to reach the target states in all of them fully.

The remainder of this paper is structured as follows: Section II presents related works. In Section III, the UX evaluation methodology is discussed. Section IV lists the UX improvements made to OMA-ML. The results from the new UX evaluation are presented in Section V. Finally, Section VI concludes the paper and discusses future works.

II. RELATED WORK

AI systems are gaining increasing importance, with new powerful AI applications being released regularly. Most recently, AI systems using Generative AI have gained prominence with applications such as ChatGPT [9], offering *human-like interactions* but also enabling new ways of powering AI systems such as code completion tools [10]. While the underlying AI algorithm is an important aspect of an AI system’s success, it also depends on how it interacts with the user [11]. This is why UX is important, as it represents a collection of strategies for understanding a user’s needs and behaviours with the system to create useful, stable systems and services [12].

However, in the past, the AI, UX and Human Computer Interaction (HCI) communities applied AI and specifically ML on a more technical approach for the creation of new methods to support the UX process itself [13] or the creation of new interfaces to interact with systems (e.g. voice interfaces) [14] [15]. The HCI research community also proposes guidelines for Human-AI Interactions [16]. In the AI community, a focus is emerging for a user-centred approach for AI systems. One prominent research area in AI that applies this is Explainable AI (XAI) [17]. XAI is a research field that emerged to focus on explaining the decision-making of AI models to the user [18] and providing insight into the data [19]. Understanding how a model came about a decision can increase usability and give the user confidence in the system, and usability is a critical part of the UX. Usability assesses how easy a user interface is to use and refers to methods for improving the ease of use during the system’s design process [20]. Our previous work proposed a list of recommended measures to ensure a good

¹<https://github.com/hochschule-darmstadt/MetaAutoML>

²<https://www.iso.org/obp/ui/#iso:std:iso:9241:-110:ed-2:v:1:en>

UX for AI systems [6]. These recommendations are based on the interaction principles formulated in ISO 9241-110:2020³ and the results from the usability study of OMA-ML. The UX methodology used for this study is introduced in the next section.

III. METHODOLOGY

Evaluating the usability of an AI platform can be a three-stage process. First, a *Use Case Model* [21] is developed to set the scope of the UX research. The use cases are extracted from the user interaction concept for this research. Next, an *Expert Review* is conducted. During the Expert Review, a UX expert inspects a system to uncover usability issues. Expert Reviews assess the design by heuristics and guidelines or principles [22]. An important set of guidelines to assess usability is the ISO Standard 9241-110:2020⁴, specifically the 7 outlined interaction principles: suitability for the user's tasks, self-descriptiveness, conformity with user expectations, learnability, controllability, use error robustness, user engagement. These interaction principles are used to measure the AI platform's usability. Based on the Use Case Model, the UX expert determines the target state for each interaction principle. The interaction principles are measured on a scale of 1 to 5 [6]. The UX Expert then evaluates the AI platform to uncover usability issues and rate the AI platform's actual state for each interaction principle. To rate the state, the German ISO 9241-110 [23] provides a checklist to determine whether the interaction principle requirements have been met. Expert Review is an effective method for catching issues. However, it may miss domain-specific issues or needs that would otherwise be found by the target audience [22]. This is resolved by performing a usability study. For this study, qualitative research techniques were chosen to comprehend what the users value the most in their experiences [24]. In the first usability study, a total of 8 usability tests were performed. Each Usability test consisted of *pretesting questions*, *follow-up questions* and *post-testing questions*. During the pretesting questions, general information about the participants was gathered, such as their demographics, motives, beliefs, expectations, existing approaches and prior experiences with AI platforms. Afterwards, the participants were given tasks they had to resolve using the AI platform. During this, follow-up questions were asked to uncover the motivations and expectations of their behaviour with the platform. During the interaction with the platform, the participants were invited to think aloud, sharing their way of thinking. Finally, after completing the work tasks, the participants were asked post-testing questions. These questions aimed to gather their feedback on the overall user experience (For a more in-depth explanation of the used usability study process, see [6]). Throughout the usability test, usability issues were collected from the participants and compiled into a list containing 104 UX issues. Each issue was categorised by the interaction principle that it infringes and

given a severity rating. This severity indicates how urgently the issue must be addressed. It is based on the five severity levels by Nielsen Norman Group [25] (0: not a usability problem, 1: cosmetic problem, 2: minor usability problem, 3: major usability problem, 4: usability catastrophe). Afterwards, OMA-ML was updated resolving each UX issue. The major UX improvements are introduced in the next section.

Finally, the entire UX evaluation was repeated with the updated version of OMA-ML and the same methodology. First, an Expert Review was conducted, collecting new issues and determining the usability state of the updated OMA-ML platform. The new state was compared to the target state and the baseline, which are the first Expert Review results using a radar chart. Then, a usability study with 29 new participants was performed. The presentation of the UX evaluation results can be seen in Section V.

IV. UX IMPROVEMENTS

In the first usability study of OMA-ML, a total of 104 UX issues were recorded. Some of the UX issues were minor, such as misunderstandings of button functionalities due to ambiguous icons or labels. For example, the dataset upload button depicted a cloud icon. This led to confusion with some participants, as they believed the dataset would be uploaded into a cloud service. In fact, most issues were major usability problems related to the participants having issues understanding what to do on a page or with elements on a page.

We identified three problems with OMA-ML which required a rework for better usability: (A) First-time user onboarding: when participants used the platform for the first time, they were unsure how to proceed or what to do; (B) Self-descriptiveness: Participants from both user groups had difficulties understanding what some of the displayed information meant or what they were supposed to do; (C) Explainable AI: The information provided by the XAI modules were too convoluted that even AI Experts did not understand what they were looking at and quickly lost interest.

To address the first two problems, we followed the 10 usability heuristics [26], most importantly, heuristic number 10: *Help and Documentation* by implementing different types of help systems. Two types of help systems can be used to help a user: *Proactive Help*, and *Reactive Help* [27]. The goal of Proactive Help is to help the user familiarize with an a user interface. This can be achieved by one of two revelations: (1) *Push Revelations*: The application provides help context without regard to the user's task, (2) *Pull Revelations*: the applications provide contextual information to the user's task. The second help system type is *Reactive Help*, which aims to answer questions and troubleshoot problems [27].

An AI system may provide a better UX if both help system types are present. Within OMA-ML, this is accomplished by providing an *interactive walkthrough* and contextual help using *tooltips* for proactive help, as well as a *documentation and search* pages for reactive help.

³<https://www.iso.org/obp/ui/#iso:std:iso:9241:-110:ed-2:v1:en>

⁴<https://www.iso.org/obp/ui/#iso:std:iso:9241:-110:ed-2:v1:en>

The XAI problem was addressed by reworking the modules. The existing modules were replaced by packages developed by the data science community. While these packages do not advertise with a focus on usability, they are popular based on their GitHub stars rating. Having an understandable XAI module is imperative for any AI platform. It helps the user understand their data and ML models. Moreover, it may increase the trust in the AI platform [28]. Making the AI platform more transparent and providing an understandable explanation is important for adopting the AI platform [29].

In the following part, the individual components used to improve the UX of OMA-ML are presented.

Interactive Walkthrough

An interactive walkthrough is a technique used for more complex applications to facilitate onboarding for new users. Onboarding is the process during which users get familiar with a new interface [30]. While it is recommended to let users experience the application independently and that tutorials such as a walkthrough may have no positive impact [31][32], they could be helpful in the context of complex AI systems, specifically AI platforms such as OMA-ML. An interactive walkthrough may ease the onboarding, enabling them to learn by doing [30]. In Fig. 1, a screenshot from the OMA-ML home dashboard page with the enabled interactive walkthrough can be seen. The current walkthrough step explains the card *Recent datasets* and instructs the user to select a dataset to proceed.

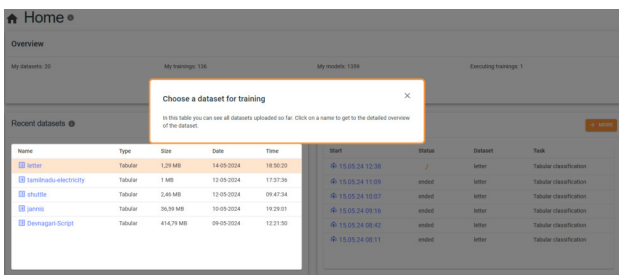


Fig. 1. OMA-ML home dashboard page with enabled interactive walkthrough

The interactive walkthrough greets first-time users upon their first login and should be like a practice run of the AI system. The OMA-ML walkthrough covers the user interaction concept [6]. At any point during the walkthrough, the user can prematurely exit and explore the platform independently. However, the platform offers the option on the documentation page to restart the walkthrough whenever the user wishes.

Documentation and Search

Documentation is an important part of UX. The main goal of documentation for reactive help is to help with user questions, troubleshoot their problems, and provide further detailed documentation for users aspiring to become expert users [27]. To achieve this, the documentation should follow some guidelines [27]: (A) It must be comprehensive and detailed; (B) it should be written following the rules of the web [33]; (C) it should make use of graphics and videos as

secondary information source; (D) optimize for search; (E) group help topics into relevant categories; (F) highlight top content that is frequently viewed.

In Fig 2, a screenshot from the documentation page of OMA-ML can be viewed.

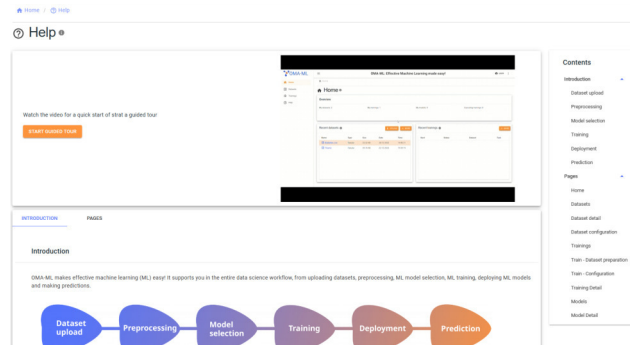


Fig. 2. OMA-ML documentation page

The documentation page consists of two sections. In the upper section, the user has a button to restart the interactive walkthrough and can view an explanatory video that follows the user interaction concept [6] and explains the process and individual pages.

The lower section provides graphical and brief written documentation for the individual steps of the user interaction concept [6] and each page within OMA-ML. Quick access links are available for the user on the right of the documentation page, listing the process steps and individual pages.

Furthermore, search functionality is available, as shown in Fig. 3. The search function aggregates all the knowledge from the underlying Ontology (See [34] for more information) and the documentation page.

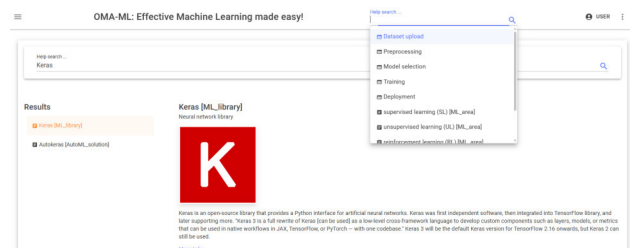


Fig. 3. OMA-ML search page

A user can search any keyword used within the platform and is then presented with a search result page. Depending on the search result, a written explanation with a graphic is shown or the corresponding help page section is displayed. In this example, the search word was *Keras* and the result are *Keras* the ML library⁵ and the AutoML solution *AutoKeras*⁶.

⁵<https://keras.io/>

⁶<https://autokeras.com/>

Tooltip

A tooltip is a brief, informative message that appears when a user interacts with an element in a graphical user interface (GUI) [35]. Tooltips are one method that can be used as a Pull Revelation for the user, providing information at the moment it is needed [32]. However, it is important to respect guidelines when incorporating them in an AI system [35][32]. Most importantly, they shall not be used to provide vital information for the user to complete their task and be used consistently. This was one of the major UX issues in the first OMA-ML study. Business domain experts and AI experts had difficulties understanding the meaning of the AI terminology in the context of the AI platform because there was no explanation. While business domain experts would not have the general background expertise to understand the meanings, AI experts would also question their expectations. For example, the word *training* in the context of data science refers to the process of training a ML model. In the context of OMA-ML, this references the Meta AutoML process of managing the training process of different AutoML solutions. This was addressed by including tooltips for any AI-specific term displayed in the platform. Depending on the nature of the element, one of two different approaches was used. First, buttons and selection options display a tooltip by hovering over them. For example, in Fig. 4, the tooltip briefly describes the selectable option of tabular classification.




Fig. 4. OMA-ML tooltip help for tabular classification

Secondly, a popup tips element was used for any element displaying information or requesting input. A popup tip is the sister element of the tooltip normally used for touchscreen devices [35]. It is paired with an "i" icon instead of being paired with an element. The OMA-ML example can be seen in Fig. 5. In this screenshot, the mouse hovers over the information icon next to a domain-specific term within the platform.

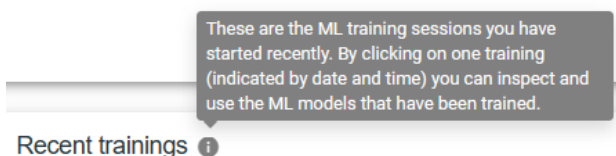


Fig. 5. OMA-ML tooltip help for recent trainings

In this case, the tooltip briefly explains what the card *Recent training* displays and what happens when clicking on one element within that card. This is a *progressive disclosure* approach [32], as it makes the existence of the tooltip visible to the user. Teaching the user information is not only available

with interaction elements but also labels or elements to provide AI knowledge as well as task explanations.

XAI

Explainable AI provides a suite of techniques that enable human users to understand, trust and produce more explainable models [19]. It is an important aspect of an AI system, almost as important as the main AI functionality, as the trust a user has towards an AI system influences the adoption decision of the AI system [29]. AI Explainability can be accomplished by incorporating techniques from the four XAI categories [19]. (A) *Data Explainability*: provides visualisation of the data giving insight into the dataset; (B) *Model Explanation*: provides techniques to understand the decision-making within black and white box models; (C) *Feature-Based Techniques*: methods to describe how input features contribute to the model output; (D) *Example-Based Techniques*: Techniques to provide explainability using dataset specific examples.

XAI research provides a toolkit of techniques to make the data and models explainable [19]. However, there is also ready-to-use modules available covering one or multiple XAI categories.

Two third-party XAI ready-to-use modules were incorporated for the XAI module of OMA-ML. The first being *ydata-profiling*⁷ for the Data Explainability. Ydata-profiling provides an Exploratory Data Analysis (EDA) by automatically performing univariate, multivariate, text, file analysis and discovers dataset challenges. In Fig. 6, a screenshot of the final EDA dashboard generated by data-profiling within OMA-ML can be seen.




Fig. 6. OMA-ML dataset analysis

In the screenshot, only a section of the dashboard can be seen; this section displays the correlation matrix between dataset features.

The second XAI module is *explainer dashboard*⁸. This XAI module generates an interactive dashboard by analysing an ML model with a corresponding dataset. It supports techniques from the remaining XAI categories (Model Explanations, Feature-Based Techniques and Example-Based Techniques). In Fig. 7, a screenshot of the explainer dashboard can be seen within OMA-ML

Depending on the ML model, different information is displayed. In this case, the dashboard contains information about the importance of features, classification stats, individual predictions, what-ifs, and feature dependence. In the screenshot,

⁷<https://github.com/ydataai/ydata-profiling>

⁸<https://github.com/oegedijk/explainerdashboard>




Fig. 7. OMA-ML explainable dashboard

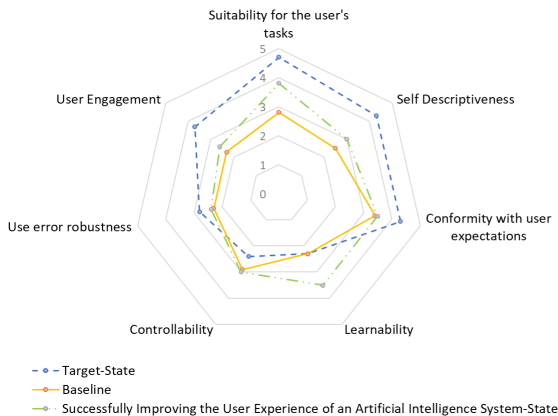


Fig. 8. Spider chart comparing the target state vs the results from the first and second Usability Study

the what-if tab is displayed. This tab provides functionality to experiment with the feature values and live evaluate how the model adjusts its prediction probability.

While neither module lists UX as a focus of their work, their popularity can be deduced from the number of GitHub stars they received (12.1k for ydata-profiling and 2.2k for explainer dashboard as of May 2024). The data science community is actively using and continuously improving these tools.

V. UX EVALUATION OF OMA-ML

Results from the Expert Review. In the first Expert Review [6], a total of four interaction principals with weak points were identified (suitability for the user's tasks, self-descriptiveness, conformity with user expectations and user engagement). Using the updated version of OMA-ML the second Expert Review could uncover that progress in the weak points could be made except for the interaction principle: conformity with user expectations. Furthermore, the learnability of OMA-ML was significantly improved. Although the learnability was already at its target goal, the new UX improvements, while aimed to improve other interaction principles, also increased the learnability. The resulting radar chart is shown in Fig. 8. The chart presents, for each of the 7 interaction principles, the different OMA-ML states on a scale of 1 to 5. The blue data points representing the target state, determined before the first Expert Review. The orange data points represent the baseline state after the first Expert Review. Finally, the green

data points show the state of the updated OMA-ML version after the current Expert Review.

Results from the Usability Study. A total of 29 usability tests were performed, collecting a total of 120 usability problems. Each usability problem was assigned the interaction principle it infringes: suitability for the user's tasks (26), self-descriptiveness (27), conformity with user expectations (38), learnability (2), controllability (7), use error robustness (14), user engagement (6). Next, a severity rating was performed using the method described in Section III and potential resolution approaches added to each usability issue. Compared to the first Usability Study, the majority (78) of the usability issues have a severity rating of 2 (minor usability problem) or lower. Some of the more notable issues were:

- 1) XAI module: when the user wants to open the XAI explainer dashboard in some cases the dashboard did not load and causes errors visible to the user;
- 2) walkthrough: in some instances, the user was unaware of the required action to proceed with the walkthrough and became stuck.

VI. CONCLUSIONS AND FUTURE WORK

UX focuses on creating a positive and meaningful experience for users and takes on a critical role during the design and development phase of any application. A bad UX can lead to rejection by the user. This is of particular importance when it comes to AI systems. As one part of UX, the usability gets determined by the trust of the underlying AI system. This is especially important to AI platforms such as OMA-ML. As without trust, the user may not adopt the platform. However, UX has not yet been a major focus in AI research. Nonetheless, the XAI research area is providing a suite of tools to support AI systems to provide explainability and trust in AI systems.

For the purpose of understanding and successfully improving the UX of an AI system, it is important to take an *iterative* qualitative and user-centric research approach. Using Expert Review and Usability tests based on the 7 interaction principles. The Expert Review uncovers UX issues, and the Usability Tests help to find usability problems from the target user group's perspective. Afterwards, the issues are resolved, and the updated version should be re-evaluated to uncover further UX issues and weak points. Using the case study of OMA-ML, we could show that the UX improvements to the interface and new XAI modules improved the platform's UX in 3 of the 4 interaction principle weak points previously uncovered after the initial UX evaluation.

While OMA-ML is not yet reaching its target state in all interaction principles, further UX improvements may be seen after the next UX evaluation iteration. The Usability study found a total of 120 new UX issues. After resolving these issues, a new iteration of the UX evaluation can be performed, potentially uncovering new ways to improve OMA-ML. Performing iterative UX evaluations per our recommendations can lead to successfully improving the UX of an AI system, potentially leading to more human-centered AI.

ACKNOWLEDGMENT

This work is funded by the German federal ministry of education and research (BMBF) in the program Zukunft der Wertschöpfung (funding code 02L19C157), and supported by Projektträger Karlsruhe (PTKA). The responsibility for the content of this publication lies with the authors.

REFERENCES

- [1] S. J. Russell and P. Norvig, *Artificial intelligence: A modern approach*, fourth edition ed., ser. Pearson Series in Artificial Intelligence. Hoboken, NJ: Pearson, 2021. ISBN 9780134610993
- [2] B. Ríos-Sánchez, D. C.-d. Silva, N. Martín-Yuste, and C. Sánchez-Ávila, "Deep learning for face recognition on mobile devices," *IET Biometrics*, vol. 9, no. 3, pp. 109–117, 2020. doi: 10.1049/iet-bmt.2019.0093
- [3] P. Polakova and B. Klimova, "Using deepl translator in learning english as an applied foreign language - an empirical pilot study," *Heliyon*, vol. 9, no. 8, p. e18595, 2023. doi: 10.1016/j.heliyon.2023.e18595
- [4] D. A. Norman, *The design of everyday things*, revised and expanded editions ed. Cambridge, MA and London: The MIT Press, 2013. ISBN 9780262525671
- [5] M. Mahmoud, U. Badawi, W. Hassan, Y. Alomari, F. Alghamdi, and T. Farag, "Evaluation of user experience in mobile applications," vol. 15, p. 2021, 2021.
- [6] L. Brand, B. G. Humm, A. Krajewski, and A. Zender, "Towards improved user experience for artificial intelligence systems," in *Engineering applications of neural networks*, ser. Communications in Computer and Information Science, S. Alonso, L. Iliadis, C. Jayne, I. Maglogiannis, and E. Pimenidis, Eds. Cham: Springer, 2023, vol. 1826, pp. 33–44. ISBN 978-3-031-34203-5
- [7] B. G. Humm, H. Bense, M. Fuchs, B. Gernhardt, M. Hemmje, T. Hoppe, L. Kaupp, S. Lothary, K.-U. Schäfer, B. Thull, T. Vogel, and R. Wenning, "Machine intelligence today: applications, methodology, and technology," *Informatik Spektrum*, pp. 1–11, 2021. doi: 10.1007/s00287-021-01343-1. [Online]. Available: <https://link.springer.com/article/10.1007%2Fs00287-021-01343-1>
- [8] A. Zender, B. G. Humm, and T. Pachmann, "Improving the efficiency of meta automl via rule-based training strategies," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F708 pp. 235–246.
- [9] P. P. Ray, "Chatgpt: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope," *Internet of Things and Cyber-Physical Systems*, vol. 3, pp. 121–154, 2023. doi: 10.1016/j.iotcps.2023.04.003
- [10] A. Sarkar, "Will code remain a relevant user interface for end-user programming with generative ai models?" in *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, T. van der Storm and R. Hirschfeld, Eds. New York, NY, USA: ACM, 2023. doi: 10.1145/3622758.3622882. ISBN 9798400703881 pp. 153–167.
- [11] V. S. Moustakis and J. Herrmann, "Where do machine learning and human-computer interaction meet?" *Applied Artificial Intelligence*, vol. 11, no. 7–8, pp. 595–609, 1997. doi: 10.1080/088395197117948
- [12] B. Pennington, S. Chapman, A. Fry, A. Deschenes, and C. G. McDonald, "Strategies to improve the user experience," *Serials Review*, vol. 42, no. 1, pp. 47–58, 2016. doi: 10.1080/00987913.2016.1140614
- [13] Å. Stige, E. D. Zamani, P. Mikalef, and Y. Zhu, "Artificial intelligence (ai) for user experience (ux) design: a systematic literature review and future research agenda," *Information Technology & People*, 2023. doi: 10.1108/ITP-07-2022-0519
- [14] M. Chromik, F. Lachner, and A. Butz, "MI for ux? - an inventory and predictions on the use of machine learning techniques for ux research," in *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, ser. ACM Digital Library, D. Lamas, H. Sarapuu, I. Šmorgun, and G. Berget, Eds. New York, NY, United States: Association for Computing Machinery, 2020. doi: 10.1145/3419249.3420163. ISBN 9781450375795 pp. 1–11.
- [15] A. Abdul, J. Vermeulen, D. Wang, B. Y. Lim, and M. Kankanhalli, "Trends and trajectories for explainable, accountable and intelligible systems," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. ACM Conferences, R. Mandryk, Ed. New York, NY: ACM, 2018. doi: 10.1145/3173574.3174156. ISBN 9781450356206 pp. 1–18.
- [16] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collislon, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, and E. Horvitz, "Guidelines for human-ai interaction," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. ACM Digital Library, S. Brewster, Ed. New York, NY, United States: Association for Computing Machinery, 2019. doi: 10.1145/3290605.3300233. ISBN 9781450359702 pp. 1–13.
- [17] D. Wang, Q. Yang, A. Abdul, and B. Y. Lim, "Designing theory-driven user-centric explainable ai," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. ACM Digital Library, S. Brewster, Ed. New York, NY, United States: Association for Computing Machinery, 2019. doi: 10.1145/3290605.3300831. ISBN 9781450359702 pp. 1–15.
- [18] A. Colley, M. Kalving, J. Häkkinä, and K. Väänänen, "Exploring tangible explainable ai (tangxai): A user study of two xai approaches," in *Bowen, Pantidi et al. (Hg.) 12.02.2023 – Proceedings of the 35th Australian*, pp. 679–683.
- [19] R. Dwivedi, D. Dave, H. Naik, S. Singhal, R. Omer, P. Patel, B. Qian, Z. Wen, T. Shah, G. Morgan, and R. Ranjan, "Explainable ai (xai): Core ideas, techniques, and solutions," *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–33, 2023. doi: 10.1145/3561048
- [20] Jakob Nielsen, "Usability 101: Introduction to usability." [Online]. Available: <https://www.nngroup.com/articles/usability-101-introduction-to-usability/>
- [21] K. Bittner and I. Spence, *Use case modeling*, ser. The Addison-Wesley object technology series. Boston, Mass.: Addison-Wesley, 2003. ISBN 0201709139
- [22] Aurora Harley, "Ux expert reviews." [Online]. Available: <https://www.nngroup.com/articles/ux-expert-reviews/>
- [23] "Din en iso 9241-110:2020-10, ergonomie der mensch-system-interaktion_ teil_110: Interaktionsprinzipien (iso_9241-110:2020); deutsche fassung en_iso_9241-110:2020," Berlin.
- [24] S. B. Merriam, *Qualitative research and case study applications in education*, rev. and expanded. ed., ser. A joint publication of the Jossey-Bass education series and the Jossey-Bass higher and adult education series. San Francisco, Calif.: Jossey-Bass Publishers, 1998. ISBN 0787910090
- [25] Jakob Nielsen, "Severity ratings for usability problems." [Online]. Available: <https://www.nngroup.com/articles/how-to-rate-the-severity-of-usability-problems/>
- [26] —, "10 usability heuristics for user interface design." [Online]. Available: <https://www.nngroup.com/articles/ten-usability-heuristics/>
- [27] Alita Joyce, "Help and documentation (usability heuristic #10)." [Online]. Available: <https://www.nngroup.com/articles/help-and-documentation/>
- [28] J. J. Ferreira and M. S. Monteiro, "What are people doing about xai user experience? a survey on ai explainability research and practice," in *Design, User Experience, and Usability. Design for Contemporary Interactive Environments*, ser. Springer eBook Collection, A. Marcus and E. Rosenzweig, Eds. Cham: Springer International Publishing and Imprint Springer, 2020, vol. 12201, pp. 56–73. ISBN 978-3-030-49759-0
- [29] Meg Kurdziolek, "Explaining the unexplainable: Explainable ai (xai) for ux." [Online]. Available: <https://uxpamagazine.org/explaining-the-unexplainable-explainable-ai-xai-for-ux/>
- [30] Alita Joyce, "Mobile-app onboarding: An analysis of components and techniques." [Online]. Available: <https://www.nngroup.com/articles/mobile-app-onboarding/>
- [31] —, "Mobile tutorials: Wasted effort or efficiency boost?" [Online]. Available: <https://www.nngroup.com/articles/mobile-tutorials/>
- [32] Page Laubheimer, "Onboarding tutorials vs. contextual help." [Online]. Available: <https://www.nngroup.com/articles/onboarding-tutorials/>
- [33] Amy Schade, "Inverted pyramid: Writing for comprehension." [Online]. Available: <https://www.nngroup.com/articles/inverted-pyramid/>
- [34] A. Zender and B. G. Humm, "Ontology-based meta automl," *Integrated Computer-Aided Engineering*, vol. 29, no. 4, pp. 351–366, 2022. doi: 10.3233/ICA-220684
- [35] Alita Joyce, "Tooltip guidelines." [Online]. Available: <https://www.nngroup.com/articles/tooltip-guidelines/>

Task-driven single-image super-resolution reconstruction of document scans

Maciej Zyrek, Michal Kawulok

0009-0009-4709-2743, 0000-0002-3669-5110

Department of Algorithmics and Software, Silesian University of Technology

Akademicka 16, 44-100 Gliwice, Poland

Email: prozyr@gmail.com; michal.kawulok@polsl.pl

Abstract—Super-resolution reconstruction is aimed at generating images of high spatial resolution from low-resolution observations. State-of-the-art super-resolution techniques underpinned with deep learning allow for obtaining results of outstanding visual quality, but it is seldom verified whether they constitute a valuable source for specific computer vision applications. In this paper, we investigate the possibility of employing super-resolution as a preprocessing step to improve optical character recognition from document scans. To achieve that, we propose to train deep networks for single-image super-resolution in a task-driven way to make them better adapted for the purpose of text detection. As problems limited to a specific task are heavily ill-posed, we introduce a multi-task loss function that embraces components related with text detection coupled with those guided by image similarity. The obtained results reported in this paper are encouraging and they constitute an important step towards real-world super-resolution of document images.

I. INTRODUCTION

INSUFFICIENT image spatial resolution is often a bottleneck for computer vision systems that limits the capabilities of image analysis algorithms. In order to address that obstacle, considerable research attention has been paid to developing techniques for image enhancement [1] and super-resolution (SR) [2] aimed at reconstructing high-resolution (HR) images from low-resolution (LR) observations, being either a single image [3] or multiple images presenting the same scene [4].

Potentially, SR algorithms can be extremely valuable in the cases when acquiring an HR image is subject to a trade-off with the acquisition cost (e.g., in remote sensing [5]), speed (e.g., in document scanning [6]), or other factors [7]. However, the attempts to apply SR algorithms as a preprocessing step prior to fulfilling a proper image analysis task are still rather scarce—commonly, the techniques are trained and validated relying on HR reference images, which are downsampled and degraded to simulate the input LR images. As noted in an excellent review by Chen et al. [3], deep networks trained from the simulated data render overoptimistic results and their performance in real-world conditions is much worse, when they are applied to enhancing original, rather than downsampled images. There have been some attempts reported to address this problem relying on the use of real-world data for training [8], [9], but acquiring such data is challenging

and costly, and it is not straightforward to exploit the HR references when HR and LR images are captured using different sensors [10]. Another possibility to regularize the training performed from the simulated data is to combine the low-level computer vision task of SR reconstruction with high-level ones like semantic segmentation [11], object detection [12], [13] and recognition [14], [15]. However, this research direction has not been extensively explored so far.

A. Related Work

Existing SR techniques can be roughly categorized into single-image (SISR) [3] and multi-image (MISR) [4] ones. The latter also embrace methods specialized for video [16] and burst-image SR [17]. While MISR techniques underpinned with information fusion are more successful in recovering the actual HR information, they are also definitely more challenging to apply, as multiple images of the same scene must be acquired and co-registered at subpixel precision. As these restrictions turn out to be impractical in many real-life cases, SISR techniques can be straightforwardly applied and they received much larger research attention. With the advent of deep learning, the field of SISR experienced unprecedented advancements [18] which nowadays allow for generating realistic images even at large magnification ratios of $8\times$ and more [19]. The first convolutional neural network (CNN) for SR (SR-CNN) [20] already outperformed the techniques based on sparse coding, despite a relatively simple architecture, which was extended and accelerated to create a faster FSRCNN [21]. The subsequent advancements adopted the achievements in feature representation and nonlinear mapping to modeling the relation between LR and HR images [22]. The larger models included a very deep SR (VDSR) network [23], deep Laplacian pyramid network (LapSRN) with progressive upsampling [24], enhanced deep SR network (EDSR) [25], and SRResNet with residual connections [26] which was used as a generator in a generative adversarial network (GAN) setting. The latest trends in SISR are more focused on reducing the size of the deep models, while preserving the reconstruction quality [27]. Recently, it was demonstrated that SISR can benefit from vision transformers [28] which dynamically adjust the size of the feature maps, thus reducing the model complexity.

There have been also some reported attempts to employ SISR to improve text detection and optical character recog-

This work was supported by the National Science Centre, Poland, under Research Grant 2022/47/B/ST6/03009.

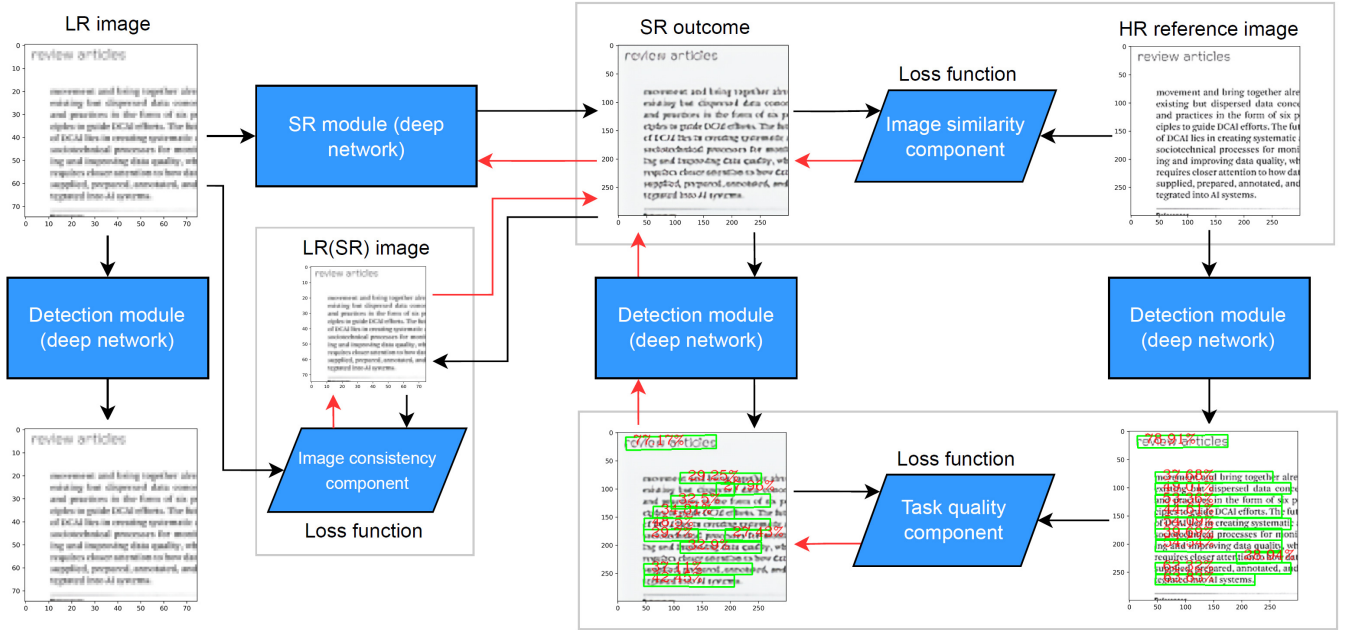


Fig. 1. Outline of the proposed self-supervised task-driven training underpinned with text detection. Red arrows indicate the propagation of the loss functions, and the black arrows show the data flow.

tion (OCR). Dong et al. adapted their SRCNN for that purpose [29] and in [30] a network with a fairly simple architecture with three convolutional layers was employed for super-resolving document scans. Wang et al. proposed to enrich a GAN-based approach with text perceptual loss to help the generator produce recognition-friendly information [31] and later they introduced a TextZoom dataset [6] composed of cropped texts from the RealSR dataset [32] with natural images captured in uncontrolled environment. In [33], a text-focused SR method was introduced which employs a vision transformer to extract sequential information. Inspired by the Gestalt psychology, stroke-based text priors were proposed in [34] and text priors were exploited for training an SR network in [35].

The aforementioned SR techniques were trained to enhance images for OCR, relying on loss functions that are correlated with that specific computer vision task. In addition to that, it is also possible to train an SR network in a task-driven manner, in which the task itself is exploited as a loss function to optimize the network's parameters. Haris et al. applied an object detection loss [12] which although leads to worse peak signal-to-noise ratio (PSNR) scores than relying on the image-similarity L1 loss, but object detection from the super-resolved images is much more effective. Similar task-driven loss functions were also defined for semantic image segmentation [36], [37]. However, in all these cases the ground-truth references related with the specific task are required for the training data.

B. Contribution

In this paper, we report our work on task-driven SISR aimed at improving text detection for an OCR system. Our

contribution can be summarized as follows:

- 1) We propose a multi-task training underpinned with a loss function composed of image-similarity and text detection-based components.
- 2) The individual components of the loss function are dynamically balanced during the network training to ensure that all components are optimized at a similar pace, even though they have different magnitudes and learning speeds.
- 3) We propose a self-supervised approach to task-driven training, with the reference labels automatically extracted from the HR reference images.
- 4) We report the results of an extensive experimental study which demonstrates that the proposed technique enhances text detection accuracy from document scans super-resolved using three different SR methods.

II. PROPOSED APPROACH

The proposed task-driven training scheme is outlined in Fig. 1. An SR network is trained using three types of loss functions: (i) similarity with the HR reference image, (ii) consistency component—similarity between the downsampled SR outcome and the input LR image, and (iii) task quality components—the similarity in the space of deep features extracted using a network that performs text detection and recognition. For image-based loss components, we employ the L2 metric (they are termed L2-HR and L2-LR for reference-based and consistency components, respectively), while for computing the task-based loss, we rely on the L1 distance.

In our study, we employed¹ the connectionist text proposal network (CTPN) for text detection [38]. For CTPN, we exploit 512 features from the final fully-connected layer (we term them as CTPN-deep), as well as the final outputs that encode the coordinates and confidence scores (20 features each, termed CTPN-out). During training, we compute the distances for these three feature spaces and we treat them as different tasks in our setup. For training SR networks, we used a CTPN model that has been already trained—its parameters are frozen during task-driven training and the gradient is propagated to optimize the SR network. Importantly, we establish the target text positions based on the outcome of text detection in the HR reference images. In this way, we do not need the text positions to be annotated, making the training self-supervised.

Our initial attempts to exploit a loss function constructed from multiple components revealed that it is quite challenging to ensure the stability between them during training. Even if we weigh these components to provide a proper initial balance, the training is becoming focused on those that are easier to be optimized and the problem turns into an imbalanced one over time. In order to address that issue, we employed a dynamic weight averaging (DWA) algorithm that adjusts the weights assigned to the particular tasks based on their individual improvements observed in subsequent training steps [39]. In DWA, for N tasks, the weight assigned to an x -th task at t -th step is determined as:

$$w_x(t) = N \exp \frac{r_x(t-1)}{T} \bigg/ \sum_{i=1}^N \exp \frac{r_i(t-1)}{T}, \quad (1)$$

where

$$r_i(t) = L_i(t)/L_i(t-1). \quad (2)$$

L_i is the value of the i -th loss component and T is the temperature controlling the softness of the task weighting (here, $T = 1$). In this way, the larger weights are assigned to these tasks in the t -th step whose losses decreases less in the preceding $(t-1)$ -th step. This makes the training more focused on these tasks that are more difficult to optimize and it prevents a single component from dominating the training process or being neglected.

III. EXPERIMENTS

In our experiments, we exploited three types of datasets: (i) natural MS COCO images [40] for training baseline SR models, (ii) scans from the benchmark datasets: Old Books² and LRDE Document Binarization Dataset (LRDE-DBD)³ [41], and (iii) our *scanned documents* dataset with real-world scans performed using a Canon LiDE 400 scanner. In our study, we investigated the SRCNN [20], FSRCNN [21] and SRResNet [26] techniques for SR at $4\times$ magnification factor. We selected these networks, as they are easy to train, while having a different level of architecture complexity. For

training these methods using the regular image-based loss function (L2-HR), we exploited the MS COCO images (LR images were obtained by downsampling the HR images) and for task-based training, we exploited a training set extracted from the Old Books and LRDE-DBD datasets (70% images). The test sets were formed from the remaining 30% of Old Books and LRDE-DBD datasets, as well as from all the scanned documents (we used five different scans split into 864 patches with 512×512 pixels). The CTPN model was trained beforehand from the ICDAR2017 dataset [42] and its parameters were frozen during the task-driven trainings.

The reconstruction quality was measured relying on image similarity metrics, namely PSNR, structural similarity index (SSIM), and learned perceptual image patch similarity (LPIPS) [43], computed between the super-resolved image and the HR reference (thus, reflecting the L2-HR loss function). For assessing the text detection quality, we employed intersection over union (IoU) between the text positions detected in the super-resolved image and in the corresponding HR reference. We also report the distances in the CTPN-deep and CTPN-out feature spaces that are used for computing the task-based components of the loss function.

First, we trained each network from scratch (60 epochs), guiding the training using a standard baseline configuration (with the L2-HR loss) and using all loss components, including L2-HR, the consistency (L2-LR) and task-based CTPN-deep and CTPN-out components. For FSRCNN and SRResNet, we fine-tuned the baseline models (100 epochs) relying on (i) L2-HR loss combined with the task-based loss components, (ii) the task component coupled with the consistency loss, and (iii) using all loss components. In addition to that, we trained SRResNet (as the best performing model) from scratch relying only on the task-based components (hence without using the image similarity at all). In Table I, we report the scores obtained for two test sets (unseen during training): for the test set of the benchmark datasets and for our dataset with the scanned documents. It can be observed that incorporating the task-based components improves the scores in terms of the image-based metrics in most cases and it always improves the quality of the text detection task (the differences are definitely higher for our scanned documents). It is also clear that the models cannot be trained from scratch without using the image-based components—apparently the problem is not convex enough and the training gets stuck in a local minimum.

A sample of the qualitative results is presented in Fig. 2 for a benchmark image and one of our scans (the configurations presented in the figure are referenced from Table I). While for the benchmark image (two upper rows), the text quality is consistently good across all configurations, for our scan, it is definitely better for the model fine-tuned in a task-driven way (d), and it is actually quite close to the result obtained in the HR reference. It can also be seen that the texts are quite clear when SRResNet is trained without using the image-based loss components (which also leads to good detection outcome), but the stability in the color space is not preserved, leading to extremely poor quantitative scores reported earlier in Table I.

¹For CTPN, we use implementation available at <https://github.com/courao/ocr.pytorch>

²Available at <https://github.com/PedroBarcha/old-books-dataset>

³Available at <https://www.lrde.epita.fr/wiki/Olena/DatasetDBD>

TABLE I
 QUANTITATIVE SCORES OBTAINED FOR THE IMAGES FROM THE OLD BOOKS AND LRDE-DBD BENCHMARKS AND FROM OUR DATASET WITH DOCUMENT SCANS, OBTAINED USING DIFFERENT SR TECHNIQUES TRAINED WITH A VARIETY OF LOSS FUNCTIONS. FOR EACH METRIC AND CATEGORY, THE BEST RESULT IS BOLDFACED.

Model and training type (a reference in Fig. 2)	Loss function				Image similarity metrics			Text detection metrics		
	L2-HR	L2-LR	CTPN-deep	CTPN-out	PSNR↑ (dB)	SSIM↑	LPIPS↓	IoU↑	CTPN-deep↓ ($\cdot 10^{-2}$)	CTPN-out↓ ($\cdot 10^{-2}$)
Test set from the simulated benchmark images (Old Books and LRDE-DBD):										
SRCNN (from scratch) (a)	✓	✓	✓	✓	21.16	0.8481	0.1818	0.8923	—	—
SRCNN (from scratch)	✓	✓	✓	✓	21.08	0.8489	0.1897	0.9290	1.831	3.366
FSRCNN (from scratch) (b)	✓	✓	✓	✓	25.17	0.9134	0.1790	0.9332	—	—
—fine-tuned	✓	✓	✓	✓	24.00	0.9071	0.2982	0.9604	1.005	1.750
—fine-tuned	✓	✓	✓	✓	20.06	0.6394	0.4681	0.9559	0.993	1.742
—fine-tuned	✓	✓	✓	✓	24.59	0.8848	0.3471	0.9560	1.113	1.939
FSRCNN (from scratch)	✓	✓	✓	✓	24.54	0.8880	0.3245	0.9588	1.097	1.919
SRResNet (from scratch) (c)	✓	✓	✓	✓	24.10	0.9147	0.1553	0.9392	—	—
—fine-tuned	✓	✓	✓	✓	28.16	0.9537	0.1037	0.9614	0.676	1.177
—fine-tuned	✓	✓	✓	✓	24.67	0.8404	0.3048	0.9676	0.694	1.198
—fine-tuned	✓	✓	✓	✓	28.04	0.9578	0.0993	0.9761	0.714	1.235
SRResNet (from scratch)	✓	✓	✓	✓	25.49	0.9302	0.1614	0.9590	1.036	1.802
SRResNet (from scratch) (e)	✓	✓	✓	✓	2.97	-0.1832	0.7714	0.9197	2.564	4.737
Scanned documents:										
SRCNN (from scratch) (a)	✓	✓	✓	✓	16.83	0.5709	0.4301	0.7103	—	—
SRCNN (from scratch)	✓	✓	✓	✓	17.06	0.5782	0.4344	0.7275	2.827	5.342
FSRCNN (from scratch) (b)	✓	✓	✓	✓	18.68	0.6542	0.3681	0.7341	—	—
—fine-tuned	✓	✓	✓	✓	18.84	0.6467	0.3585	0.7608	2.363	4.290
—fine-tuned	✓	✓	✓	✓	16.39	0.4796	0.4343	0.7688	2.294	4.174
—fine-tuned	✓	✓	✓	✓	18.82	0.6429	0.3588	0.7635	2.328	4.219
FSRCNN (from scratch)	✓	✓	✓	✓	18.82	0.6444	0.3644	0.7641	2.414	4.348
SRResNet (from scratch) (c)	✓	✓	✓	✓	18.70	0.6634	0.3798	0.7264	—	—
—fine-tuned	✓	✓	✓	✓	19.62	0.7075	0.3189	0.7910	1.985	3.397
—fine-tuned	✓	✓	✓	✓	19.00	0.6327	0.3261	0.7886	1.994	3.520
—fine-tuned	✓	✓	✓	✓	19.81	0.7076	0.3164	0.7807	2.023	3.483
SRResNet (from scratch)	✓	✓	✓	✓	19.23	0.6731	0.3591	0.7576	2.361	4.280
SRResNet (from scratch) (e)	✓	✓	✓	✓	2.91	-0.0929	0.9250	0.7186	3.170	5.592

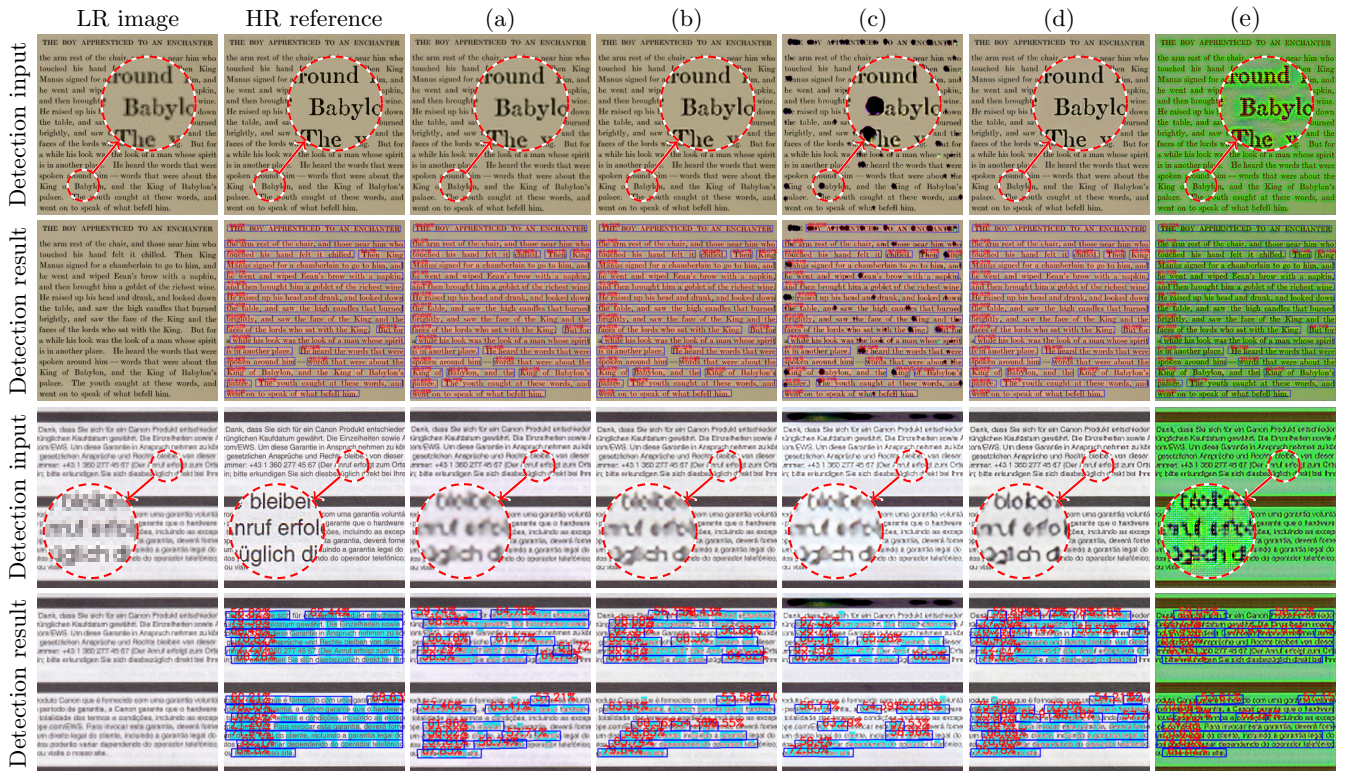


Fig. 2. Examples of SR reconstruction with: (a) SRCNN, (b) FSRCNN and (c) SRResNet (all with L2-HR loss), (d) fine-tuned SRResNet (L2-HR, L2-LR, CTPN-deep and CTPN-out loss functions), and (e) SRResNet trained from scratch (CTPN-deep and CTPN-out loss functions). These settings are also referred to in Table I. For each example (top: Old Books; bottom: our dataset), we include the detection input (i.e., the SR outcome) and the detected text.

IV. CONCLUSIONS AND OUTLOOK

In this paper, we reported our initial attempts to apply task-driven training for SISR, guided by text detection. The results are highly encouraging, revealing high potential of task-based loss functions. Importantly, in contrast to the earlier works concerned with task-driven SR, we train the models in a self-supervised way, as we retrieve the annotations by processing the HR reference images.

Our ongoing research is focused on including the text recognition components that may improve the guidance during training. Also, we plan to adapt our approach to MISR problems and to create a dataset embracing samples composed of multiple scans of the same document.

REFERENCES

- [1] F. Jelowicki, "Enhancing image quality through automated projector stacking," in *Communication Papers of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. A. Maciaszek, M. Paprzycki, and D. Slezak, Eds., vol. 37, 2023, pp. 153–156. [Online]. Available: <https://doi.org/10.15439/2023F9900>
- [2] W. Yang, X. Zhang, Y. Tian, W. Wang, J. Xue, and Q. Liao, "Deep learning for single image super-resolution: A brief review," *IEEE Transaction on Multimedia*, vol. 21, no. 12, pp. 3106–3121, Dec 2019. [Online]. Available: <https://doi.org/10.1109/TMM.2019.2919431>
- [3] H. Chen, X. He, L. Qing, Y. Wu, C. Ren, R. E. Sheriff, and C. Zhu, "Real-world single image super-resolution: A brief review," *Information Fusion*, vol. 79, pp. 124–145, 2022. [Online]. Available: <https://doi.org/10.1016/j.inffus.2021.09.005>
- [4] L. Yue, H. Shen, J. Li, Q. Yuan, H. Zhang, and L. Zhang, "Image super-resolution: The techniques, applications, and future," *Signal Processing*, vol. 128, pp. 389–408, 2016. [Online]. Available: <https://doi.org/10.1016/j.sigpro.2016.05.002>
- [5] T. Tarasiewicz, J. Nalepa, R. A. Farrugia, G. Valentino, M. Chen, J. A. Briffa, and M. Kawulok, "Multitemporal and multispectral data fusion for super-resolution of Sentinel-2 images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–19, 2023. [Online]. Available: <https://doi.org/10.1109/TGRS.2023.3311622>
- [6] W. Wang, E. Xie, X. Liu, W. Wang, D. Liang, C. Shen, and X. Bai, "Scene text image super-resolution in the wild," in *Proc. IEEE/CVF ECCV*. Springer, 2020, pp. 650–666. [Online]. Available: https://doi.org/10.1007/978-3-030-58607-2_38
- [7] T. Balon, M. Knapik, and B. Cyganek, "Real-time detection of small objects in automotive thermal images with modern deep neural architectures," in *Communication Papers of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. A. Maciaszek, M. Paprzycki, and D. Slezak, Eds., vol. 37, 2023, pp. 29–35. [Online]. Available: <https://doi.org/10.15439/2023F8409>
- [8] J. Cai, S. Gu, R. Timofte, and L. Zhang, "NTIRE 2019 Challenge on real image super-resolution: Methods and results," in *Proc. IEEE/CVF CVPR*, 2019, pp. 1–13. [Online]. Available: <https://doi.org/10.1109/CVPRW.2019.00274>
- [9] M. Märten, D. Izzo, A. Krzic, and D. Cox, "Super-resolution of PROBA-V images using convolutional neural networks," *Astrodynamics*, vol. 3, no. 4, pp. 387–402, 2019. [Online]. Available: <https://doi.org/10.1007/s42064-019-0059-8>
- [10] P. Kowaleczko, T. Tarasiewicz, M. Ziaja, D. Kostrzewa, J. Nalepa, P. Rokita, and M. Kawulok, "A real-world benchmark for Sentinel-2 multi-image super-resolution," *Scientific Data*, vol. 10, no. 1, p. 644, 2023. [Online]. Available: <https://doi.org/10.1038/s41597-023-02538-9>
- [11] Z. Guo, G. Wu, X. Song, W. Yuan, Q. Chen, H. Zhang, X. Shi, M. Xu, Y. Xu, R. Shibasaki *et al.*, "Super-resolution integrated building semantic segmentation for multi-source remote sensing imagery," *IEEE Access*, vol. 7, pp. 99 381–99 397, 2019. [Online]. Available: <https://doi.org/10.1109/ACCESS.2019.2928646>
- [12] M. Haris, G. Shakhnarovich, and N. Ukita, "Task-driven super resolution: Object detection in low-resolution images," in *Proc. ICONIP*. Springer, 2021, pp. 387–395. [Online]. Available: https://doi.org/10.1007/978-3-030-92307-5_45
- [13] T. Balon, M. Knapik, and B. Cyganek, "New thermal automotive dataset for object detection," in *Position Papers of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. A. Maciaszek, M. Paprzycki, and D. Slezak, Eds., vol. 31, 2022, pp. 43–48. [Online]. Available: <https://doi.org/10.15439/2022F283>
- [14] X. Yang, W. Wu, K. Liu, P. W. Kim, A. K. Sangaiah, and G. Jeon, "Long-distance object recognition with image super resolution: A comparative study," *IEEE Access*, vol. 6, pp. 13 429–13 438, 2018. [Online]. Available: <https://doi.org/10.1109/ACCESS.2018.2799861>
- [15] M. Włodarczyk-Sielicka and D. Polap, "Interpolation merge as augmentation technique in the problem of ship classification," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. A. Maciaszek, and M. Paprzycki, Eds., vol. 21, 2020, pp. 443–446. [Online]. Available: <https://doi.org/10.15439/2020F11>
- [16] H. Liu, Z. Ruan, P. Zhao, C. Dong, F. Shang, Y. Liu, L. Yang, and R. Timofte, "Video super-resolution based on deep learning: a comprehensive survey," *Artificial Intelligence Review*, pp. 1–55, 2022. [Online]. Available: <https://doi.org/10.1007/s10462-022-10147-y>
- [17] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte, "Deep burst super-resolution," in *Proc. IEEE/CVF CVPR*, 2021, pp. 9209–9218. [Online]. Available: <https://doi.org/10.1109/CVPR46437.2021.00909>
- [18] Z. Wang, J. Chen, and S. C. H. Hoi, "Deep learning for image super-resolution: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 10, pp. 3365–3387, 2021. [Online]. Available: <https://doi.org/10.1109/TPAMI.2020.2982166>
- [19] R. Abiantun, F. Juefei-Xu, U. Prabhu, and M. Savvides, "SSR2: Sparse signal recovery for single-image super-resolution on faces with extreme low resolutions," *Pattern Recognition*, vol. 90, pp. 308–324, 2019. [Online]. Available: <https://doi.org/10.1016/j.patcog.2019.01.032>
- [20] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Proc. IEEE/CVF ECCV*. Springer, 2014, pp. 184–199. [Online]. Available: https://doi.org/10.1007/978-3-319-10593-2_13
- [21] C. Dong, C. C. Loy, and X. Tang, "Accelerating the super-resolution convolutional neural network," in *Proc. IEEE/CVF ECCV*. Springer, 2016, pp. 391–407. [Online]. Available: https://doi.org/10.1007/978-3-319-46475-6_25
- [22] Y. Huang, J. Li, X. Gao, Y. Hu, and W. Lu, "Interpretable detail-fidelity attention network for single image super-resolution," *IEEE Transactions on Image Processing*, vol. 30, pp. 2325–2339, 2021. [Online]. Available: <https://doi.org/10.1109/TIP.2021.3050856>
- [23] J. Kim, J. Kwon Lee, and K. Mu Lee, "Accurate image super-resolution using very deep convolutional networks," in *Proc. IEEE/CVF CVPR*, 2016, pp. 1646–1654. [Online]. Available: <https://doi.org/10.1109/CVPR.2016.182>
- [24] W. Lai, J. Huang, N. Ahuja, and M. Yang, "Fast and accurate image super-resolution with deep Laplacian pyramid networks," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 41, no. 11, pp. 2599–2613, 2019. [Online]. Available: <https://doi.org/10.1109/TPAMI.2018.2865304>
- [25] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE/CVF CVPR Workshops*, 2017, pp. 136–144. [Online]. Available: <https://doi.org/10.1109/CVPRW.2017.151>
- [26] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE/CVF CVPR*, 2017, pp. 4681–4690. [Online]. Available: <https://doi.org/10.1109/CVPR.2017.19>
- [27] M. Ayazoglu, "Extremely lightweight quantization robust real-time single-image super resolution for mobile devices," in *Proc. IEEE/CVF CVPR*, 2021, pp. 2472–2479. [Online]. Available: <https://doi.org/10.1109/CVPRW53098.2021.00280>
- [28] Z. Lu, J. Li, H. Liu, C. Huang, L. Zhang, and T. Zeng, "Transformer for single image super-resolution," in *Proc. IEEE/CVF CVPR*, 2022, pp. 457–466. [Online]. Available: <https://doi.org/10.1109/CVPRW56347.2022.00061>

- [29] C. Dong, X. Zhu, Y. Deng, C. C. Loy, and Y. Qiao, "Boosting optical character recognition: A super-resolution approach," *arXiv preprint arXiv:1506.02211*, 2015. [Online]. Available: <https://doi.org/10.48550/arXiv.1506.02211>
- [30] R. K. Pandey and A. Ramakrishnan, "Efficient document-image super-resolution using convolutional neural network," *Sādhanā*, vol. 43, pp. 1–6, 2018. [Online]. Available: <https://doi.org/10.1007/s12046-018-0794-1>
- [31] W. Wang, E. Xie, P. Sun, W. Wang, L. Tian, C. Shen, and P. Luo, "TextSR: Content-aware text super-resolution guided by recognition," *arXiv preprint arXiv:1909.07113*, 2019. [Online]. Available: <https://doi.org/10.48550/arXiv.1909.07113>
- [32] J. Cai, H. Zeng, H. Yong, Z. Cao, and L. Zhang, "Toward real-world single image super-resolution: A new benchmark and a new model," in *Proc. IEEE ICCV*, 2019. [Online]. Available: <https://doi.org/10.1109/ICCV.2019.00318>
- [33] J. Chen, B. Li, and X. Xue, "Scene text telescope: Text-focused scene image super-resolution," in *Proc. IEEE/CVF CVPR*, 2021, pp. 12 026–12 035. [Online]. Available: <https://doi.org/10.1109/CVPR46437.2021.01185>
- [34] J. Chen, H. Yu, J. Ma, B. Li, and X. Xue, "Text Gestalt: Stroke-aware scene text image super-resolution," in *Proc. AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 285–293. [Online]. Available: <https://doi.org/10.1609/aaai.v36i1.19904>
- [35] J. Ma, S. Guo, and L. Zhang, "Text prior guided scene text image super-resolution," *IEEE Transactions on Image Processing*, vol. 32, pp. 1341–1353, 2023. [Online]. Available: <https://doi.org/10.1109/TIP.2023.3237002>
- [36] T. Frizza, D. G. Dansereau, N. M. Seresht, and M. Bewley, "Semantically accurate super-resolution generative adversarial networks," *Computer Vision and Image Understanding*, p. 103464, 2022. [Online]. Available: <https://doi.org/10.1016/j.cviu.2022.103464>
- [37] M. S. Rad, B. Bozorgtabar, C. Musat, U.-V. Marti, M. Basler, H. K. Ekenel, and J.-P. Thiran, "Benefiting from multitask learning to improve single image super-resolution," *Neurocomputing*, vol. 398, pp. 304–313, 2020. [Online]. Available: <https://doi.org/10.1016/j.neucom.2019.07.107>
- [38] Z. Tian, W. Huang, T. He, P. He, and Y. Qiao, "Detecting text in natural image with connectionist text proposal network," in *Proc. IEEE/CVF ECCV*. Springer, 2016, pp. 56–72. [Online]. Available: https://doi.org/10.1007/978-3-319-46484-8_4
- [39] S. Vandenhende, S. Georgoulis, W. Van Gansbeke, M. Proesmans, D. Dai, and L. Van Gool, "Multi-task learning for dense prediction tasks: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3614–3633, 2021. [Online]. Available: <https://doi.org/10.1109/TPAMI.2021.3054719>
- [40] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Proc. IEEE/CVF ECCV*. Springer, 2014, pp. 740–755. [Online]. Available: https://doi.org/10.1007/978-3-319-10602-1_48
- [41] G. Lazzara and T. Géraud, "Efficient multiscale sauvola's binarization," *International Journal on Document Analysis and Recognition (IJDAR)*, vol. 17, no. 2, pp. 105–123, 2014. [Online]. Available: <https://doi.org/10.1007/s10032-013-0209-0>
- [42] R. Gomez, B. Shi, L. Gomez, L. Numann, A. Veit, J. Matas, S. Belongie, and D. Karatzas, "ICDAR2017 robust reading challenge on COCO-text," in *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, vol. 1. IEEE, 2017, pp. 1435–1443. [Online]. Available: <https://doi.org/10.1109/ICDAR.2017.234>
- [43] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE/CVF CVPR*, 2018. [Online]. Available: <https://doi.org/10.1109/CVPR.2018.00068>

Congestion Control in Streaming Services with an On-Off MPTCP Algorithm

Łukasz Piotr Łuczak
0000-0003-0892-7276

Institute of Information Technology
Lodz University of Technology, Łódź, Poland
Email: lukasz.luczak@dokt.p.lodz.pl

Przemysław Ignaciuk
0000-0003-4420-9941

Institute of Information Technology
Lodz University of Technology, Łódź, Poland
Email: przemyslaw.ignaciuk@p.lodz.pl

Abstract—In this paper, by adopting the analytical framework of dynamical systems theory, a new congestion control (CC) algorithm for Multipath TCP (MPTCP) streaming services is developed. The proposed nonlinear algorithm, following an on-off principle, is formally demonstrated robust with respect to variable network conditions. It maintains stream consistency at the negotiated video rate despite *a priori* unknown delay fluctuations. Also, precise guidelines for buffer size allocation at the receiver are provided. The designed algorithm is compared against the established MPTCP CC algorithms: LIA, OLIA, BaLIA, and wVegas. The tests, conducted in open networks using real devices and equipment, show that the on-off controller reduces protocol delay, jitter, and Head-of-Line blocking, which is crucial for ensuring high-quality streaming in mobile networks.

Index Terms—MPTCP, congestion control, streaming applications, tactile Internet.

I. INTRODUCTION

THE RAPID expansion of Internet Protocol (IP)-based systems across a spectrum of applications underscores the need for network solutions that are both flexible and scalable. Despite their pervasive use, IP networks encounter dynamic resource allocation challenges that compromise Quality of Service (QoS), particularly affecting streaming services that demand a consistent, high-quality data transfer rate [1], [2]. By engaging multiple network paths, Multipath TCP (MPTCP) has been recognized as a promising solution to improve transmission reliability and efficiency. The core of MPTCP's effectiveness lies in its congestion control (CC) algorithms, which are crucial for managing data flow across multiple paths. However, the core MPTCP CC algorithms, namely LIA, OLIA, BaLIA, and wVegas, do not address the particular requirements of streaming applications in terms of minimizing delay and jitter [3].

Recent research in the development of CC algorithms for MPTCP has led to the proposal of several strategies aimed at improving network throughput and user fairness while efficiently managing network resources. A CC algorithm proposed by Kou et al. relies on packet loss to optimize network throughput and fairness among users, presenting an

approach to balance resource allocation [4]. Similarly, the D-LIA algorithm, developed by Lubna et al., dynamically adjusts the congestion window decrease factor based on packet loss intervals, targeting throughput enhancement and packet loss reduction [5]. Additionally, Mudassir and Baig introduced the Modified Fast-Vegas-LIA Hybrid Congestion Control Algorithm (MFVL HCCA) for MPTCP, which operates in both uncoupled and coupled congestion control modes to adapt to network conditions, resulting in significant improvements in packet loss reduction and average goodput [6]. These contributions illustrate ongoing efforts to refine MPTCP CC algorithms to meet the requirements of modern network applications.

Our team has also contributed to this field through various publications. For instance, we explored the performance of MPTCP in industrial Internet of Things (IoT) applications and proposed a green multipath TCP framework to enhance energy efficiency and network performance [7]. Additionally, we conducted an experimental assessment of MPTCP congestion control algorithms specifically for streaming services in open Internet environments, highlighting the performance differences among LIA, OLIA, BaLIA, and wVegas algorithms [8]. In another study, we investigated appropriate control strategies for multipath transmission in Industry 4.0 applications, emphasizing the need for adaptive congestion control mechanisms [9].

This study introduces an On-Off CC MPTCP algorithm, designed to enhance streaming performance by adequately responding to the dynamic nature of network conditions. Unlike conventional approaches, the design of the On-Off algorithm, by adopting the analytical framework of dynamic systems theory, allows for a formal analysis of its properties. Following a rigorous mathematical argument, the proposed algorithm is demonstrated to be robust to variable networking conditions manifesting themselves in delay fluctuations. The receiver's data queue is shown to be finite with a precisely estimated upper limit that constitutes the required buffer capacity. Moreover, conditions for maintaining the preestablished transfer rate, and thus video quality, are formulated and proved.

The evaluation of the On-Off algorithm in a public network environment underscores its capacity to optimize congestion management, thereby directly addressing the limitations of existing MPTCP CC algorithms in maintaining high-quality

This work has been performed in the framework of a project "Robust control solutions for multi-channel networked flows" no. 2021/41/B/ST7/00108 financed by the National Science Centre, Poland

streaming services under fluctuating network conditions. Traditional MPTCP CC mechanisms, while proficient in handling multiple paths, were not specifically designed for scenarios where rapid adjustments to bandwidth fluctuations are essential to preserve streaming quality. This inadequacy manifests as increased latency and jitter, which are detrimental to real-time applications. The On-Off algorithm introduces a more dynamic and responsive approach to congestion control by design, effectively minimizing these adverse effects and enhancing the landscape of MPTCP CC solutions. This advancement promises improvements in the stability and reliability of streaming services, marking an enhancement over predecessors such as LIA, OLIA, BaLIA, and wVegas, which do not cater specifically to the high-demand scenarios of modern streaming applications [10], [11].

II. RELATED WORK

To place this work in context, we review recent contributions in the area of MPTCP congestion control algorithms and highlight previous works by our team.

A. Multipath TCP Congestion Control Algorithms

Recent studies have made significant advancements in MPTCP congestion control. For instance, Kou et al. proposed a congestion control algorithm based on packet loss, which optimizes network throughput and fairness among users by balancing resource allocation [4]. Similarly, the D-LIA algorithm by Lubna et al. dynamically adjusts the congestion window decrease factor based on packet loss intervals, enhancing throughput and reducing packet loss [5]. Mudassar and Baig's MFVL HCCA algorithm adapts to network conditions by operating in both uncoupled and coupled modes, resulting in improved packet loss reduction and average goodput [6]. Additional research has explored other novel approaches such as a BBR-based congestion control scheme aimed at improving throughput in heterogeneous wireless networks [12], and a delay-based congestion control algorithm designed to optimize transmission performance by minimizing delay differences between paths [13]. The advanced MPTCP (AMP) protocol adjusts congestion detection based on subflow count, enhancing latency for small flows and throughput for large flows [14]. Another notable approach is the Shared Bottleneck-based Congestion Control scheme (SB-CC) utilizing ECN to detect shared bottlenecks among subflows, improving network performance [15].

B. Our Contributions

Our team has also contributed to this field through various publications. For instance, we explored the performance of MPTCP in industrial Internet of Things (IoT) applications and proposed a green multipath TCP framework to enhance energy efficiency and network performance [7]. Additionally, we conducted an experimental assessment of MPTCP congestion control algorithms specifically for streaming services in open Internet environments, highlighting the performance

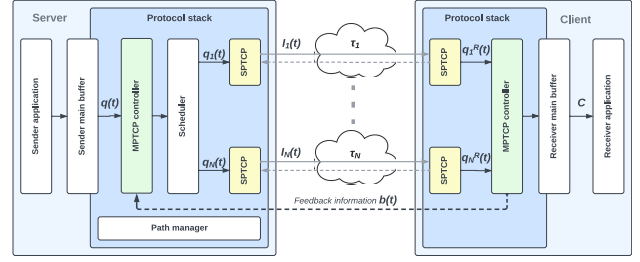


Fig. 1: Client-server interaction in an MPTCP framework.

differences among LIA, OLIA, BaLIA, and wVegas algorithms [8]. In another study, we investigated appropriate control strategies for multipath transmission in Industry 4.0 applications, emphasizing the need for adaptive congestion control mechanisms [9]. These works underscore our ongoing efforts to refine MPTCP congestion control to meet the evolving requirements of modern network applications.

Furthermore, the proposed On-Off CC MPTCP algorithm builds on our previous work by integrating dynamic systems theory to offer a more robust solution for streaming applications. This approach ensures better adaptability to fluctuating network conditions, thereby enhancing the quality and reliability of streaming services.

III. SYSTEM MODEL

A. System Variables

Consider the scenario depicted in Fig. 1, where a user needs uninterrupted access to a live video broadcast at the negotiated rate C while on the move. To enhance reliability and counteract variable network conditions, the system employs independent connections. This multi-connection configuration is enabled by MPTCP, a protocol that enhances Internet connectivity by allowing the simultaneous use of multiple links.

In this model, the intervals $t = 0, 1, 2, \dots$ represent discrete time points for video transmission. The function $s_i(t)$ denotes the number of segments transmitted through subpath i at time t . These segments experience varying delays $\tau_i(t)$ as they traverse the network. The number of successfully delivered segments is expressed as

$$s_i^R(t) = s_i(t - \tau_i(t)), \quad (1)$$

and the queue length in the receiving buffer evolves according to:

$$b(t+1) = b(t) + \sum_{i=1}^N s_i(t - \tau_i(t)) - b_0. \quad (2)$$

Assuming zero initial conditions, i.e., an empty receiver buffer and no in-flight data at $t = 0$, the accumulation of segments in the receiver buffer for any $t > 0$ evolves as:

$$b(t) = \sum_{i=1}^N \sum_{k=0}^{t-1} s_i(k - \tau_i(k)) - Ct. \quad (3)$$

B. Uncertainty Model

The delay $\tau_i(t)$ varies randomly within the range:

$$(1 - \delta_i)T_i \leq \tau_i(t) \leq (1 + \delta_i)T_i, \quad (4)$$

where T_i represents the nominal delay on subpath i , and $\delta_i \in [0, 1]$ indicates the uncertainty tolerance. The delay variability reflects inherent network fluctuations, such as traffic load or routing changes.

To quantify the impact of delay variations on data delivery, we introduce a function $\eta(t)$:

$$\eta(t) = \eta_+(t) - \eta_-(t), \quad (5)$$

where $\eta_+(t)$ accounts for segments that arrive sooner than expected, and $\eta_-(t)$ for those delayed beyond the expected timeframe. These components are calculated as:

$$\eta_+(t) = \sum_{i=1}^N \sum_{j \in (0, \delta T_i): \tau_i(t+j) \leq T_i - j} s_i(t - T_i + j) \quad (6)$$

representing the accumulation of segments that outpaced the anticipated delay, while

$$\eta_-(t) = \sum_{i=1}^N \sum_{j \in (0, \delta T_i): \tau_i(t-j) > T_i + j} s_i(t - T_i + j) \quad (7)$$

captures those delayed beyond the expected timeframe.

Thus, the receiver queue length evolution is given by

$$b(t) = \sum_{i=1}^N \sum_{k=0}^{t-1} s_i(k - T_i) + \eta(t) - Ct. \quad (8)$$

Considering the maximum amount of data that can be injected into path i in a time unit S_i , the magnitude of $\eta(t)$ does not exceed a maximum value, H , estimated as:

$$H = \sum_{i=1}^N \delta_i T_i S_i. \quad (9)$$

C. Congestion Control Algorithm

After evaluating the network's available resources and in-flight data, the sender prepares a set of S data segments for the transmission across N subpaths. The paths are allocated on a dynamic basis according to:

$$S_i(t) = \lambda_i(t)S, \quad (10)$$

where $S_i(t)$ is the amount of data forwarded onto path i , and $\lambda_i(t) \in [0, 1]$, reflects the allocation strategy. A complete allocation is assumed, i.e.,

$$\sum_{i=1}^N \lambda_i(t) = 1. \quad (11)$$

The default MPTCP scheduler allocates the bandwidth evenly across all the subpaths, i.e., $\lambda_i = \frac{1}{N}$.

In the proposed On-Off algorithm, the decision to transmit S data segments at time t is guided by the following criteria:

$$s(t) = \begin{cases} S, & \text{if } B_{ref} - b(t) - I(t) > 0 \\ 0, & \text{if } B_{ref} - b(t) - I(t) \leq 0 \end{cases}, \quad (12)$$

where, B_{ref} denotes a queue reference level, indicating a desirable buffer fill level, and $I(t)$ represents the amount of in-flight data, i.e., the data that have been sent but not yet acknowledged. $I(t)$ can be calculated as:

$$\begin{aligned} I(t) &= \sum_{k=0}^{t-1} s(k) - \sum_{i=1}^N \sum_{k=0}^{t-1} s_i^R(k) \\ &= \sum_{k=0}^{t-1} s(k) - \sum_{i=1}^N \sum_{k=0}^{t-1} s_i(k - \tau_i(k)), \end{aligned} \quad (13)$$

This formula encapsulates the aggregate of all dispatched segments that are pending acknowledgment, providing a real-time snapshot of in-flight data.

IV. ON-OFF ALGORITHM FORMAL ANALYSIS

The properties of the proposed algorithm are articulated through two theorems, each rigorously proven. The first theorem demonstrates that, despite the randomness of delays, the data accumulation in the receiver's buffer is bounded by a specific limit. This limit sets the required buffer capacity to effectively manage incoming data and prevent drops. The second theorem outlines how to set the reference queue length to prevent data starvation and sustain the desired streaming rate C .

For the proofs, we introduce the concept of the channel uncertainty function, which measures the actual volume of in-flight data compared to the anticipated amount:

$$\begin{aligned} I(t) &= \sum_{k=0}^{t-1} s(k) - \sum_{i=1}^N \sum_{k=0}^{t-1} s_i(k - T_i) - \eta(t) \\ &= \sum_{k=0}^{t-1} s(k) - \sum_{i=1}^N \sum_{k=0}^{t-T_i-1} s_i(k) - \eta(t). \end{aligned} \quad (14)$$

The uncertainty function calculates the volume of in-flight data by considering the total segments sent, subtracting those successfully delivered (accounting for network delays and variability), and adjusting for early or late arrivals via the $\eta(t)$ term.

Theorem 1. *Applying Algorithm (12) to system (2) results in a finite queue length at the receiver:*

$$\forall_{t \geq 0} b(t) \leq B, \quad (15)$$

with the upper bound

$$B = B_{ref} + S + H. \quad (16)$$

Proof. Initially, the receiver buffer is empty, and remains so for any $t \leq t_{\min} = \min\{(1 - \delta_i)T_i\}$. For any time $p > t_{\min}$, two cases can be distinguished:

Case 1: When $B_{ref} > s(p) + I(p)$, following from equation (14) we have:

$$\begin{aligned} b(p) &< B_{ref} - \sum_{k=0}^{p-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p-T_i-1} s_i(k) + \eta(p) \\ &= B_{ref} - \sum_{i=1}^N \sum_{k=p-T_i}^{p-1} s_i(k) + \eta(p) \leq B_{ref} + \eta(p). \end{aligned} \quad (17)$$

Given that $\eta(t)$ does not exceed its maximum expected value H for any $t \geq 0$, it follows that $b(p) < B_{ref} + H$, establishing the upper bound in this case.

Case 2: In the case where $B_{ref} \leq s(p) + I(p)$, we identify the most recent interval p_1 before p where I was less than $B_{ref} - b(p_1)$. Given the initial conditions, such an interval p_1 exists. The buffer level $b(p)$ during p_1 adheres to an inequality similar to equation (17), i.e.,

$$b(p) < B_{ref} - \sum_{k=0}^{p_1-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p_1-1} s_i(k - \tau_i(k)). \quad (18)$$

The data volume $b(p)$ can be expressed in terms of $b(p_1)$ as:

$$b(p) = b(p_1) + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k - \tau_i(k)) - C(p - 1 - p_1). \quad (19)$$

Considering that C is nonnegative, by applying (18) to (19), we deduce:

$$\begin{aligned} b(p) &< B_{ref} - \sum_{k=0}^{p_1-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p_1-1} s_i(k - \tau_i(k)) \\ &\quad + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k - \tau_i(k)) \\ &= B_{ref} + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k) + \eta(p) \end{aligned} \quad (20)$$

Given the data was last sent before p at p_1 , and taking into account the uncertainty bound H ,

$$b(p) \leq B_{ref} + S + H. \quad (21)$$

□

Theorem 2. When algorithm (12) is applied to system (2), and S exceeds C , then the reference queue length:

$$B_{ref} > S\Lambda + S + H, \quad (22)$$

where

$$\Lambda = \max_t \sum_{i=1}^N \sum_{k=t-L_i}^{t-1} \lambda_i(k), \quad (23)$$

ensures that for any $t \geq t_{\max} = \max_i \{(1+\delta_i)T_i\} + B/(S-C)$, the receiver queue length is strictly positive.

Proof. Consider a period $p \geq t_{\max}$. Two scenarios are distinguished: when $I(p) < B_{ref} - b(p)$, and when $I(p) \geq B_{ref} - b(p)$.

Case 1: In the scenario where $I(p) \geq B_{ref} - b(p)$, following from (14), it is evident that:

$$\begin{aligned} b(p) &\geq B_{ref} - \sum_{k=0}^{p-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p-T_i-1} s_i(k) + \eta(p) \\ &= B_{ref} - \sum_{i=1}^N \sum_{k=p-T_i}^{p-1} s_i(k) + \eta(p) \\ &= B_{ref} - \sum_{i=1}^N \sum_{k=p-T_i}^{p-1} \lambda_i(k)s(k) + \eta(p). \end{aligned} \quad (24)$$

The volume of data that can be dispatched at any time within the interval $[p - T_i, p - 1]$ equals S , yet,

$$\sum_{i=1}^N \sum_{k=p-T_i}^{p-1} \lambda_i(k)s(k) \leq S \sum_{i=1}^N \sum_{k=p-T_i}^{p-1} \lambda_i(k) \leq S\Lambda. \quad (25)$$

Given that $\eta \geq -H$, it follows from equation (24) that:

$$b(p) \geq B_{ref} - S\Lambda - H. \quad (26)$$

Using assumption (25), we deduce that $b(p) > 0$, thus concluding the first part of the proof.

Case 2: Consider the scenario where the in-flight data volume $I(p) < B_{ref} - b(p)$. Identify the most recent time $p_1 < p$ when I was greater than or equal to $B_{ref} - b(p)$. According to Theorem 1, the data volume at the receiver does not exceed B . Given the buffer depletion rate C , the maximum duration for continuous data reception is $B/(S - C)$, confirming the existence of period p_1 . It is inferred from the theorem's conditions that by period p_1 , the initial data packets from all channels have been received, despite any delay variations.

The volume of in-flight data $I(p_1) \geq B_{ref} - b(p_1)$, and by employing reasoning similar to equations (24) we conclude:

$$b(p) \geq B_{ref} - \sum_{k=0}^{p_1-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p_1-1} s_i(k - \tau_i(k)) > 0. \quad (27)$$

Thus,

$$\begin{aligned} b(p) &= b(p_1) + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k - \tau_i(k)) - C(p - 1 - p_1) \\ &\geq B_{ref} - \sum_{k=0}^{p_1-1} s(k) + \sum_{i=1}^N \sum_{k=0}^{p_1-1} s_i(k - \tau_i(k)) \\ &\quad + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k - \tau_i(k)) - C(p - 1 - p_1) \\ &= B_{ref} + \sum_{i=1}^N \sum_{k=p_1}^{p-1} s_i(k) - \sum_{i=1}^N \sum_{k=p-L_i}^{p-1} s_i(k) + \eta(p) \\ &\quad - C(p - 1 - p_1). \end{aligned} \quad (28)$$

TABLE I: Performance metrics of MPTCP CC algorithms

		LIA	OLIA	BALIA	wVegas	On-Off
Protocol jitter [ms]	θ_{av}	38.45	38.53	41.11	28.13	<u>15.45</u>
	θ_{max}	329.32	256.80	241.83	227.68	<u>123.87</u>
Protocol delay [ms]	v_{av}	324.82	266.53	277.23	151.80	<u>106.26</u>
	v_{max}	588.04	480.88	560.29	418.02	<u>225.54</u>
HoL Degree [ms]	ζ_{av}	251.89	195.19	205.48	83.86	<u>36.66</u>
	ζ_{max}	518.86	410.46	391.52	350.40	<u>157.20</u>
SRTT path 1 [ms]	τ_{av}^1	2.17	2.14	1.85	<u>1.72</u>	1.87
	τ_{max}^1	3.86	4.31	3.39	<u>2.30</u>	4.07
SRTT path 2 [ms]	τ_{av}^2	72.93	71.33	71.75	67.94	69.60
	τ_{max}^2	111.98	136.83	120.29	<u>100.49</u>	111.51
Mean drop rate [seg/s]	δ^1	2.90	2.90	2.90	2.90	2.90
	δ^2	1.53	1.90	<u>0.37</u>	0.50	0.57
Throughput [Mbps]	φ_{av}	5.07	5.07	<u>5.12</u>	2.57	4.90
	φ_{max}	<u>8.48</u>	8.09	7.85	5.49	7.16

At time p_1 , $I(p_1) + b(p_1) \geq B_{ref}$, marking the last instance when no new data request is made. Subsequently, a volume S of data is dispatched to the receiver. Given the complete allocation (11), the first sum in equation (28) equates to $S(p - 1 - p_1)$. Following a logic similar to (25), the second sum is approximated as:

$$\sum_{i=1}^N \sum_{k=p-T_i}^{p-1} s_i(k) = \sum_{i=1}^N \sum_{k=p-T_i}^{p-1} \lambda_i(k) s(k) \leq S\Lambda. \quad (29)$$

The buffer occupancy level described in equation (28) fulfills the inequality:

$$b(p) > B_{ref} + S(p - 1 - p_1) - S\Lambda + \eta(p) - C(p - p_1). \quad (30)$$

This denotes the maximum data shortfall due to overly delayed data packets. Therefore, using the theorem assumption (22), it can be stated that

$$b(p) > S + S(p - 1 - p_1) - C(p - p_1) > 0. \quad (31)$$

Under the stipulated conditions, the buffer at the receiver maintains a strictly positive level, ensuring continuous data flow.

□

V. EXPERIMENTAL SETUP

The experimental setup aimed to assess the performance of various congestion control (CC) algorithms in delivering streaming content over Multipath TCP (MPTCP). This setup replicates a typical data transmission scenario where a client fetches content from a server situated in a remote data center accessible via a public IP address. Both the client and server systems operated on Linux OS version 4.19, customized to support MPTCP version 0.95. For streaming content generation, we employed VLC media player, an open-source multimedia framework, player, and server. VLC was selected due to its reliability and widespread use in streaming "Big Buck Bunny," a popular open-source animated film, ensuring our

testing environment accurately reflects real-world streaming conditions. The client setup included two communication interfaces: one connected to an LTE router through Ethernet and the other connected to the same router via Wi-Fi 802.11bgn, utilizing two distinct LTE networks. Each test session lasted for 10 seconds and was repeated 30 times to ensure the results were statistically significant.

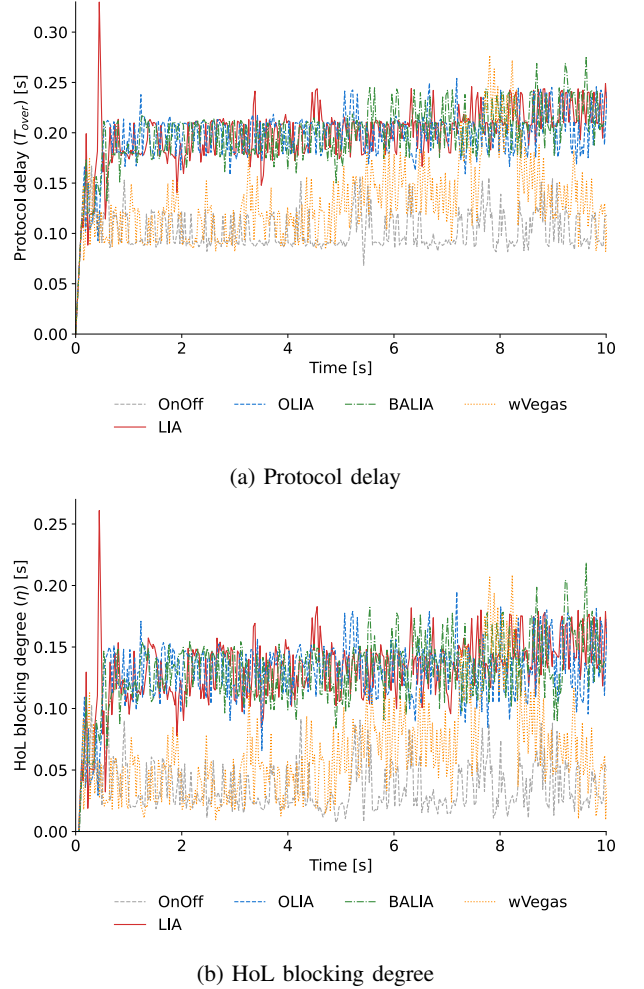


Fig. 2: Measured transmission properties.

VI. TESTS AND RESULTS

In our analysis, the On-Off algorithm was compared with several well-known CC algorithms: LIA, OLIA, BaLIA, and wVegas. These algorithms were evaluated based on metrics important for streaming applications, including protocol delay, jitter, and Head-of-Line (HoL) blocking degree. The results are summarized in Table I, and their graphical representation is shown in Fig. 2.

The On-Off algorithm demonstrated an improvement in average protocol delay, recording 106.26 ms, which is lower than the delays observed with the other algorithms. This lower delay suggests that the On-Off algorithm is more efficient in

managing the transmission of data packets, thereby reducing latency. In terms of protocol jitter, the On-Off algorithm exhibited an average jitter of 15.45 ms with a variance of 0.30, indicating less fluctuation in transmission times. This stability is important for streaming applications, as it ensures a more consistent delivery of data packets, minimizing interruptions and buffering events. Regarding HoL blocking, the On-Off algorithm achieved an average of 36.66 ms. HoL blocking can lead to delays in packet delivery, and the lower HoL blocking value for the On-Off algorithm indicates its effectiveness in maintaining a steady flow of data, which is essential for high-quality streaming.

These metrics underscore the On-Off algorithm's potential to enhance the user experience by ensuring smoother and more consistent streaming in variable network conditions. Additionally, the comparison of Smoothed Round-Trip Time (SRTT) values and drop rates further validates the algorithm's efficiency in handling multi-path data flow for seamless streaming.

VII. CONCLUSION

In this study, we evaluated the On-Off algorithm for Multipath TCP (MPTCP) congestion control (CC), aimed at enhancing streaming services via multiple network paths. Performance comparisons were made with established protocols—LIA, OLIA, BaLIA, and wVegas—highlighting key metrics such as protocol delay, jitter, and Head-of-Line (HoL) blocking. The research involved the development and analysis of a mathematical model based on dynamical systems theory, designed to predict the On-Off algorithm's performance across various network conditions. This model outlines an approach to MPTCP CC by focusing on reducing protocol delay and jitter, which are vital for delivering reliable streaming content.

Empirical tests, conducted in an Open Internet environment using real devices and equipment, validated the On-Off algorithm's effectiveness in lowering protocol delay and jitter, as well as reducing HoL blocking. These results, which align with the mathematical model's predictions, demonstrate the algorithm's capacity to enhance the streaming experience by ensuring continuous content delivery. The findings affirm the On-Off algorithm's potential to not only improve streaming quality but also to decrease receiver buffer size requirements and bolster connection stability. This positions the On-Off algorithm as an effective MPTCP congestion control strategy, particularly suitable for mobile networks.

Through both analytical and empirical evaluations, this study confirms the On-Off algorithm's utility, laying the groundwork for further research and its prospective implementation in public network applications. The On-Off algorithm's robust performance in varying network conditions highlights its potential for improving user experiences in streaming services, ensuring smoother and more reliable content delivery.

ACKNOWLEDGMENT

This work has been completed while the first author was the Doctoral Candidate in the Interdisciplinary Doctoral School at

the Lodz University of Technology, Poland

REFERENCES

- [1] "Cisco annual internet report (2018-2023) white paper," 2020. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>
- [2] K. Hatakeyama, Y. Osana, M. Tanabe, and S.-i. Kuribayashi, "Proposed congestion control method reducing the size of required resource for all-ip networks," in *2009 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*. Victoria, BC, Canada: IEEE, 8 2009. doi: 10.1109/PACRIM.2009.5291413. ISBN 978-1-4244-4560-8 pp. 1–4, [Online; accessed 2024-04-03]. [Online]. Available: <http://ieeexplore.ieee.org/document/5291413/>
- [3] C. Raiciu, M. Handley, and D. Wischik, "Coupled congestion control for multipath transport protocols," RFC Editor, Tech. Rep., 10 2011, dOI: 10.17487/rfc6356. [Online]. Available: <https://www.rfc-editor.org/info/rfc6356>
- [4] L. Kou, R. Wang, and S. R. Chen, "A mptcp congestion control algorithm based on packet loss," in *Frontiers of Manufacturing Science and Measuring Technology III*, ser. Applied Mechanics and Materials, vol. 401. Trans Tech Publications Ltd, 12 2013. doi: 10.4028/www.scientific.net/AMM.401-403.1760 pp. 1760–1765.
- [5] T. Lubna, I. Mahmud, and Y.-Z. Cho, "D-LIA: Dynamic congestion control algorithm for MPTCP," *ICT Express*, vol. 6, no. 4, pp. 263–268, 2020. doi: <https://doi.org/10.1016/j.ict.2020.03.005>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405959519304229>
- [6] M. U. Mudassir and M. I. Baig, "Mfv1 hcca: A modified fast-vegas-lia hybrid congestion control algorithm for mptcp traffic flows in multihomed smart gas iot networks," *Electronics*, vol. 10, no. 6, 2021. doi: 10.3390/electronics10060711. [Online]. Available: <https://www.mdpi.com/2079-9292/10/6/711>
- [7] M. Morawski and P. Ignaciuk, "A green multipath tcp framework for industrial internet of things applications," *Computer Networks*, vol. 187, p. 107831, 3 2021. doi: 10.1016/j.comnet.2021.107831
- [8] Ł. Łuczak, P. Ignaciuk, and M. Morawski, "Experimental assessment of mptcp congestion control algorithms for streaming services in open internet," in *Annals of Computer Science and Information Systems*, 10 2023. doi: 10.15439/2023F9991 pp. 359–364, [Online; accessed 2024-04-03]. [Online]. Available: <https://annals-csis.org/proceedings/2023/drp/9991.html>
- [9] M. Morawski and P. Ignaciuk, "Choosing a proper control strategy for multipath transmission in industry 4.0 applications," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, pp. 3609–3619, 6 2022. doi: 10.1109/TII.2021.3105499
- [10] M. O. Farooq, T. Kunz, and M. St-Hilaire, "Differentiated services based congestion control algorithm for wireless multimedia sensor networks," in *2011 IFIP Wireless Days (WD)*. Niagara Falls, ON, Canada: IEEE, 10 2011. doi: 10.1109/WD.2011.6098182. ISBN 978-1-4577-2028-4 pp. 1–6, [Online; accessed 2024-04-03]. [Online]. Available: <http://ieeexplore.ieee.org/document/6098182/>
- [11] C. Lee, "Hcca for wireless mobile coverage networks," *Indian Journal of Science and Technology*, vol. 9, no. 1, pp. 1–8, 1 2016. doi: 10.17485/ijst/2016/v9i46/107198
- [12] W. Wei, K. Xue, J. Han, Y. Xing, D. S. L. Wei, and P. Hong, "Bbr-based congestion control and packet scheduling for bottleneck fairness considered multipath tcp in heterogeneous wireless networks," *IEEE Transactions on Vehicular Technology*, vol. 70, pp. 330–345, 1 2021. doi: 10.1109/TVT.2020.3047877
- [13] H. Li, Y. Wang, R. Sun, S. Guo, and H. Wang, "Delay-based congestion control for multipath tcp in heterogeneous wireless networks," *IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, pp. 1–6, 4 2019. doi: 10.1109/WCNCW.2019.8902835
- [14] J. Ye, L. Feng, Z. Xie, J. Huang, and X. Li, "Fine-grained congestion control for multipath tcp in data center networks," *IEEE Access*, vol. 7, pp. 30495–30504, 3 2019. doi: 10.1109/ACCESS.2019.2902860
- [15] W. Wei, K. Xue, J. Han, D. S. L. Wei, and P. Hong, "Shared bottleneck-based congestion control and packet scheduling for multipath tcp," *IEEE/ACM Transactions on Networking*, vol. 28, pp. 1380–1395, 2 2020. doi: 10.1109/TNET.2020.2970032

Thematic Sessions

PARTS 4 and 5 of FedCSIS 2024 Proceedings contain contributions originating from the Thematic Sections. Let us briefly introduce each one of them (in alphabetical order).

I. ADVANCES IN PROGRAMMING LANGUAGES

Programming languages are programmers' most basic tools. With appropriate programming languages one can drastically reduce the cost of building new applications, as well as maintaining existing ones. In the last decades, there have been many advances in programming languages technology. This includes both the traditional programming paradigms, such as functional, logic, and object-oriented programming, as well as the development of new ones, such as aspect-oriented programming. The main driving force was, and will be, to better express programmers' ideas. Therefore, research in programming languages is an endless activity and the core of computer science. New language features, new programming paradigms, and better compile-time and run-time mechanisms can be foreseen in the future. Here, the future role of intelligence systems, as the target for programming, should also be considered. In this context, the aim of this thematic session was to provide a forum for exchange of ideas and experience in topics concerned with programming languages, with focus on intelligence systems development and implementation.

Thematic Session organizers:

- + Kardas, Geylani, Ege University International Computer Institute, Turkey
- + Mernik, Marjan, University of Maribor, Slovenia
- + Rangel Henriques, Pedro, Universidade do Minho, Portugal
- + Slivnik, Boštjan, University of Ljubljana, Slovenia
- + Janousek, Jan, Czech Technical University, Czech Republic
- + Varanda Pereira, Maria Joao, Instituto Politecnico de Braganca, Portugal

II. ARTIFICIAL INTELLIGENCE IN AGRICULTURE

Artificial intelligence is increasingly used in agriculture, to address multiple issues, from plant disease detection to weeding automation, soil status monitoring, crop prediction, irrigation management, and decreased use of resources, for improving product quality and process productivity. As a matter of fact, when AI methods and approaches are infused into the agriculture, the resulting Agri-intelligence Systems, can deliver, for instance, precision agriculture, by optimizing, automating and forecasting multiple aspects of farming, and revolutionizing the sector, providing helpful information and driving decisions using multiple sources of data and different sensors. Moreover, in the era of climate change, Agri-intelligence Systems may be able to improve long-term sustainability, by optimizing use of resources, e.g. in the case of

water and soil management. This thematic session welcomed contributions concerning all aspects of interdisciplinary research and applications related to Agri-intelligence Systems.

Thematic session organizers:

- + Charvat, Karel, Czech Center for Science and Society, Prague, Czech Republic
- + Martinelli, Massimo, National Research Council of Italy, Pisa, Italy
- + Moroni, Davide, National Research Council of Italy, Pisa, Italy
- + Procházka, Ales, University of Chemistry and Technology & Czech Technical University CIIRC, Prague, Czech Republic

III. ARTIFICIAL INTELLIGENCE IN DIGITAL HUMANITIES, COMPUTATIONAL SOCIAL SCIENCES AND ECONOMICS RESEARCH

This thematic session was dedicated to the computational study of social sciences, economics and humanities including, among others, education, labor market, history, religious studies, cultural heritage, and informative predictions for decision-making and behavioral-science perspectives. Besides new discoveries, it was devoted to reflections about their growth within computer science and emphasized the interdisciplinary exchange and dissemination with a clear focus on computational and intelligence-oriented methods. Since there is a clear methodological overlap between the considered domains of social sciences, economics and humanities and, often, similar algorithms and AI approaches are considered for them, this session was a place for discussing a “joint toolbox”, supporting scholars from multiple fields with human and context-aware “agents”. The session included also research related to trustworthy data infrastructures, housing both quantitative and qualitative data.

Thematic session organizers:

- + Cooper, Anthony-Paul, Durham University, Durham, United Kingdom and University of Turku, Turku, Finland
- + Dörpinghaus, Jens, BIBB and University of Koblenz, Koblenz, Germany
- + Helmrich, Robert, BIBB and University of Bonn, Bonn, Germany
- + Speckesser, Stefan, Brighton University, Brighton, United Kingdom

IV. APPLICATION OF DISRUPTIVE TECHNOLOGIES FOR SOCIETY 5.0

This thematic session provided an opportunity to discuss recent disruptive technologies and their application in the new Society 5.0, for better human life. It considered new trends in blockchain, AI and Big Data applied, among others, to fields related to smart cities and Society 5.0 such as healthcare, education, finance, mobility, logistics and quality of life. It provided an opportunity to present and discuss the

most recent and relevant research results, innovations, experiences, concerns, challenges, and trends about the application of disruptive technologies applied to Society 5.0.

Thematic session organizers:

- + Arezki, Sara, Hassan First University, Morocco
- + Khan, Inam Ullah, Isra University, Pakistan
 - + Ebersold, Sophie, University of Toulouse, IRIT CNRS, France
 - + El Hamlaoui, Mahmoud, Mohamed V University in Rabat, Morocco

V. CHALLENGES FOR NATURAL LANGUAGE PROCESSING

This thematic session consisted of contributions related to all aspects of NLP. Of particular interest were works addressing NLP tools, multimodal problems, cross-lingual learning and processing of natural languages.

Thematic session organizers:

- + Kobyliński, Łukasz, Institute of Computer Science, Polish Academy of Sciences, Poland
- + Kubis, Marek, Faculty of Mathematics and Computer Science, Adam Mickiewicz University, Poland

VI. COMPLEX NETWORKS: THEORY AND APPLICATION

In the world around us, one can observe many network structures that interconnect various elements such as cells, people, urban centers, network devices, companies, manufacturing machines, etc. Moreover, it is easy to notice that most of them evolve over time. The analysis of such systems, from the point of view of complex networks, brings about better understanding of the processes within them, which can be used to optimize their structure, improve their management methods, detect failures, improve their operating efficiency or plan their development and evolution. Currently, each of these aspects separately and all of them jointly are based on application of tools and methods from artificial intelligence. Moreover, the ultimate result of the majority of relevant work is to support development of intelligence systems. Therefore, the main goal of this thematic session was to exchange knowledge and experience between specialists from different areas who, in their research and design work, use theories and solutions characteristic for complex systems. This thematic session was organized in cooperation with Regional Center of Excellence for Automation and Robotics, Computer Science, Electrical Engineering, Electronics and Telecommunications Rzeszów University of Technology

Thematic session organizers:

- + Kondratenko, Yuriy, Petro Mohyla Black Sea National University, Ukraine
- + Paszkiewicz, Andrzej, Rzeszów University of Technology, Poland

VII. COMPUTATIONAL OPTIMIZATION

Many real-world problems, arising in engineering, economics, medicine and other domains, can be formulated as optimization tasks. These problems are frequently characterized by non-convex, non-differentiable, discontinuous, noisy or dynamic objective functions and constraints, which ask for adequate computational methods. The aim of this thematic session was to stimulate communication between researchers working in different fields of optimization, and practitioners who need reliable and efficient computational optimization methods. Contributions related to both theoretical and practical aspects of optimization methods were represented.

Thematic session organizers:

- + Fidanova, Stefka, Bulgarian Academy of Sciences, Bulgaria
- + Mucherino, Antonio, IRISA, University of Rennes, France
- + Zaharie, Daniela, West University of Timisoara, Romania

VIII. COMPUTER ASPECTS OF NUMERICAL ALGORITHMS

Numerical algorithms are widely used by scientists engaged in various areas (including computational kernels, used to realize modern machine learning algorithms). There is a special need for highly efficient and easy-to-use scalable tools for solving large scale problems. This thematic session was devoted to numerical algorithms, with particular attention focused on the latest scientific trends in this area and on problems related to implementation of libraries of efficient numerical algorithms. The main goal of this track was to facilitate meeting of researchers and exchange of their experiences.

Thematic session organizers:

- + Bylina, Beata, Maria Curie-Skłodowska University, Lublin, Poland
- + Bylina, Jarosław, Maria Curie-Skłodowska University, Lublin, Poland
- + Cyganek, Bogusław, AGH University of Science and Technology, Kraków, Poland
- + Lirkov, Ivan, Bulgarian Academy of Sciences, Sofia, Bulgaria
- + Stpicyński, Przemysław, Maria Curie-Skłodowska University, Lublin, Poland

IX. DATA SCIENCE IN HEALTH, ECOLOGY AND COMMERCE

This thematic session was a forum for exchange of ideas concerning all forms of data analysis, data economics, information systems and data-based research, focusing on the interaction of those three fields. Here, data-driven intelligence solutions can be generated by understanding complex real-world (health-related) problems, critical thinking and analytics to derive knowledge from (Big) data. The past years have

shown a forthcoming interest in innovative data technology and analytics solutions that link and utilize large amounts of data across individual digital ecosystems. Here, scenarios, in the field of health, smart cities or agriculture, merge data from various IoT devices, social media or applications, and demonstrate the great potential for gaining new insights, supporting decisions, or providing smarter services. Together with inexpensive sensors and computing power they provide foundation of a world that bases its decisions on data. However, this is only the beginning of the journey towards intelligence systems, and the pertinent methods and technologies, and the potential application fields, as well as the impact on society and economy, must be explored. This endeavor needs the knowledge of researchers from different fields applying diverse perspectives and using different methodological directions to find a way to grasp and fully understand the power and opportunities of data science. Bringing together researchers and practitioners of pertinent fields was one of focal points of this thematic session.

Thematic session organizers:

- + Bumberger, Jan, Helmholtz-Centre for Environmental Research – UFZ, Leipzig, Germany
- + Franczyk, Bogdan, University of Leipzig, Leipzig, Germany
- + Häckl, Dennis, University of Leipzig, Germany and WIG2 Institute for Health Economics and Health Service Research, Leipzig, Germany
- + Militzer-Horstmann, Carsta, WIG2 Institute for Health Economics and Health Service Research, Leipzig, Germany
- + Reinhold, Olaf, University of Leipzig / Social CRM Research Center, Leipzig, Germany

X. INDUSTRIAL CYBER-PHYSICAL SYSTEMS AND SOFTWARE AGENTS

In line with the strategic objectives of Horizon Europe and the move towards Industry 4.0, the effective management of Industrial Cyber-Physical Systems presents economic challenges and opportunities. Addressing the complex challenges of large-scale collaboration and the nuances of distributed system architectures is essential for the development of robust and efficient ICPSs. It is therefore essential to strengthen research into predictive control systems, focusing on both theoretical and practical applications to meet the demands of these sophisticated environments. With the increasing complexity of ICPSs, there is a pressing need for advanced modelling and simulation methods. This is particularly relevant in the context of Multi-Agent Systems (MAS), which benefit from the integration of, broadly understood, artificial intelligence. The convergence of the Internet of Things, Cloud computing and Edge computing technologies is crucial in this context, enabling distributed learning and improving the responsiveness of MAS in ICPS. Hence, this thematic session seeks to explore the potential for collaboration between software agents and ICPSs, using the interconnections provided by IoT, Cloud and Edge com-

puting to meet the strategic objectives of international research agendas and the principles of Industry 4.0.

Thematic session organizers

- + Alves, Filipe, DTx – Digital Transformation Colab, Portugal
- + Barbosa, José, Research Centre in Digitalization and Intelligent Robotics (CeDRI), Polytechnic Institute of Bragança, Portugal
- + Rodrigues, Nelson Ricardo, DTx – Digital Transformation Colab, Portugal

XI. INFORMATION SYSTEMS MANAGEMENT

This thematic session facilitated a forum for exchange of ideas, for practitioners and theorists working in the broad area of intelligent information systems management in organizations. It is focused on three complimentary directions: management of information systems in an organization, uses of information systems to empower managers, and information systems for sustainable development. Here, the interest encompassed all aspects of planning, organizing, resourcing, coordinating, controlling and leading the management function to ensure a smooth operation of intelligent information systems in organizations. Moreover, the contributions discussing the uses of intelligence systems and information technology to facilitate (and, possibly, fully automate) the management functions were included. Research on the influence of intelligence systems on sustainability was welcomed as well.

Thematic session organizers:

- + Bicevska, Zane, University of Latvia, Riga, Latvia
- + Chmielarz, Witold, University of Warsaw, Warsaw, Poland
- + Leyh, Christian, University of Applied Sciences, Giessen, Germany
- + Sołtysik-Piorunkiewicz, Anna, University of Economics in Katowice, Poland

XII. INTERNET OF THINGS – ENABLERS, CHALLENGES AND APPLICATIONS

The Internet of Things (IoT) is rapidly being deployed around the world. IoT applications include, among others: smart city initiatives, wearable devices aimed to real-time health monitoring, smart homes and buildings, smart vehicles, environment monitoring, intelligent border protection, logistics support. IoT assumes pervasive presence of smart things, including sensors, actuators, embedded systems, etc. Widespread connectivity, getting cheaper smart devices and an exploding demand for data suggest that IoT will continue to grow by leaps and bounds. The business models of various industries are being redesigned around IoT principles. This thematic session was focused on the IoT challenges in

networking and information management, security and privacy, logistics, situation awareness, and medical care.

Thematic track organizers:

+ Chudzikiewicz, Jan, Military University of Technology, Poland

+ Zieliński, Zbigniew, Military University of Technology, Poland

XIII. MODEL DRIVEN APPROACHES IN SYSTEM DEVELOPMENT

For many years, various approaches in system design and implementation differentiated between the specification of the system and its implementation on a particular platform. People in software industry have been using models for a precise description of systems at the appropriate abstraction level without unnecessary details. Model Driven (MD) approaches to the system development increase the importance and power of models by shifting the focus from programming to modeling activities. Models may be used as primary artifacts in constructing software, which means that software components are generated from models. Software development tools need to automate as many as possible tasks of model construction and transformation requiring the smallest amount of human interaction. The goal of the thematic session is to bring together researchers working on MD approaches, techniques and tools, as well as Domain Specific Modeling and Domain Specific Languages and applying them in the requirements engineering, information system and application development, databases, and related areas, so that they can exchange their experience, create new ideas, evaluate and improve MD approaches and spread its use. Moreover, in response to the recent trends, an additional goal is to reflect on the relations between MD approaches and development of intelligence systems.

Thematic session organizers:

+ Milašinović, Boris, University of Zagreb, Faculty of Electrical Engineering and Computing, Croatia

+ Gray, Jeff, University of Alabama, United States

+ Ristić, Sonja, University of Novi Sad, Faculty of Technical Sciences, Serbia

XIV. MULTIMEDIA APPLICATIONS AND PROCESSING

Multimedia, computer vision, graphics, and machine learning have become ubiquitous in modern information systems, creating new challenges for detection, recognition, indexing, access, search, retrieval, automated understanding, and processing, resulting in many applications based on image and signal processing, machine learning and various multimedia technologies. Recent advances in pervasive computers, networks, telecommunications, and information technology, along with the proliferation of multimedia mobile devices, have stimulated the rapid development of new generation of multimedia-focused intelligence systems. Here, the key technologies, including (but not limited to)

virtual reality, augmented reality, and computational intelligence, facilitate a multimedia revolution that significantly impacts a broad spectrum of consumer, business, healthcare, educational and governmental domains. This thematic session covered a range of AI-anchored theories, methods, algorithms, technologies, and systems for diversified and heterogeneous digital multimedia, imaging, computer graphics and machine learning areas. Moreover, it provided an opportunity for researchers and professionals to discuss present and future challenges and potential collaboration for future progress in these fields.

Thematic Session organizers:

+ Iwanowski, Marcin, Warsaw University of Technology, Poland

+ Kwaśnicka, Halina, Wrocław University of Science and Technology, Poland

+ Śluzek, Andrzej, Khalifa University, United Arab Emirates

+ Stanescu, Liana, University of Craiova, Romania

XV. RESILIENCE IN CRITICAL INFRASTRUCTURES AND SYSTEMS

With the increased digitalization and adoption of off-the-shelf information technology components, Critical Infrastructures (CI) have benefited in many ways, but the cyber-attack surface also became larger. Besides many single CI cyber risks, inter-CI systems and software also have a need for unified prevention, protection, data processing, and fast recovery. These issues are even more pronounced in the case of intelligence systems, which start to involve (semi-)autonomous decision making. In this context, the European Commission tries to tackle the resilience of CI with the integration of different strategies, indicators, and tools that rely on intelligent processing of available data, as well as regulations and policy measures. The focus of the thematic session was to explore all aspects related to making critical infrastructures, and intelligence systems in particular, more resilient.

Thematic session organizers:

+ Blanco, Jose Miguel, Polytechnic University of Madrid (UPM), Spain

+ Jovanovic, Aleksandar, European Risk and Resilience Institute (EU-VRi), Germany

+ Pasic, Aljosa, Eviden, Spain

+ Schauer, Stefan, Austrian Institute of Technology (AIT), Austria

+ Segou, Olga, Netcompany-Intrasoft, Greece

XVI. ROUGH SETS: THEORY AND APPLICATIONS

This thematic session discussed research related to the state-of-the-art and future perspectives of rough sets, considered from both a theoretical standpoint and real-world appli-

cations. Rough set theory is a versatile mathematical framework that has proven successful in AI, knowledge representation, approximate reasoning, data mining, machine learning, and pattern recognition, among other areas. The track was devoted to all the mentioned areas, with an additional emphasis on problems of modeling AI processes using rough set-based techniques. The session provided an opportunity for interdisciplinary exchange and collaboration among scientists from diverse backgrounds, including mathematics, computer science, statistics, physics, engineering, and social sciences. Moreover, it allowed staying up to date with the state-of-the-art in rough set theory and its applications, and to discuss future research directions and opportunities.

Thematic session organizers:

- + Artiemjew, Piotr, University of Warmia and Mazury in Olsztyn, Poland
- + Chelly Dagdia, Zaineb, UVSQ, Paris-Saclay, France
- + Mani, A., Indian Statistical Institute, Kolkata, India

XVII. SCALABLE COMPUTING

The world of large-scale computing continuously evolves. The most recent addition to the mix comes from numerous data streams that materialize from exploding number of cheap sensors installed “everywhere”, on the one hand, and ability to capture and study events with systematically increasing granularity, on the other. To address the needs for scaling computational and storage infrastructures, concepts like: edge, fog and dew computing emerged. Novel issues, involved in “pushing computing away from the center”, did not replace open questions that existed in the context of Grid and Cloud computing. Rather, they added new dimensions of complexity and resulted in the need of addressing scalability across more and more complex ecosystems consisting of individual sensors and micro-computers (e.g. Raspberry

PI based systems) as well as supercomputers available within the Cloud. Moreover, intelligence systems bring about questions related to machine learning scalability. Addressing research questions that arise in individual “parts” as well as across the ecosystem viewed from a holistic perspective, with scalability as the focus, was the goal of this thematic session.

Thematic session organizers:

- + Gepner, Paweł, Warsaw University of Technology, Warsaw, Poland
- + Gusev, Marjan, University Ss. Cyril and Methodius, Skopje, Macedonia
- + Petcu, Dana, West University of Timisoara, Timisoara, Romania
- + Ristov, Sashko, University of Innsbruck, Innsbruck, Austria
- + Stencel, Krzysztof, University of Warsaw, Warsaw, Poland

XVIII. SELF LEARNING AND SELF ADAPTIVE SYSTEMS

Self-learning systems consist of entities that can acquire and renew knowledge over time without hard coding. Typically, in this class of intelligence systems, functions improve by a learning process, where the system initially interacts with its users, or the surrounding environment, by attempting and observing the changes produced by its actions. This session focused on the design, implementation and exploitation of self-learning features, within an intelligent environment or some of its components.

Thematic session organizers:

- + Coronato, Antonio, Università Giustino Fortunato, Benevento, Italy
- + Naeem, Muddasar, Università Giustino Fortunato, Benevento, Italy

Virtual Power Plant Optimization Service - Benchmark of Solvers

Filipe Alves^{0000-0002-8387-391X*}, Rui Ribeiro^{0009-0002-3421-8932*}, Maria Petiz^{0009-0001-1059-0871*},
Ali Abbasi^{0000-0002-5581-1279*}, Pedro Carvalho^{0009-0002-1319-6535*}, Ricardo Faia^{0000-0002-1053-7720†},
Pedro Faria^{0000-0002-5982-8342†}, Zita Vale^{0000-0002-4560-9544†}, and Ricardo Rodrigues^{0000-0001-7986-3754*}

* DTx — Digital Transformation CoLAB, University of Minho, 4800-058 Guimarães, Portugal

Email: {filipe.alves, rui.ribeiro, maria.petiz, ali.abbasi, pedro.carvalho, ricardo.rodrigues}@dtx-colab.pt

† GECAD - Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development,
LASI - Intelligent Systems Associate Laboratory; Polytechnic of Porto; Porto, Portugal

Email: {rff, pnf, zav}@isep.ipp.pt

Abstract—This work provides a comprehensive analysis of the optimization of a Virtual Power Plant (VPP), that consider the presence of energy storage systems and controllable loads, through the benchmarking of various solvers. It delves into the development of a Mixed Integer Linear Programming (MILP) algorithm aiming at optimizing energy management and exchange within a VPP, that takes into account the operation of shift electric appliances and battery storage systems among different houses. The proposed model aims to minimize the overall electricity cost while ensuring that the energy demand of the system is met, the battery state of charge is maintained within safe operating limits, and the shift electrical appliance is scheduled. Furthermore, the experimental comparisons, the study evaluates the performance of commercial and open-source solvers in handling the complex dynamics of energy demand and supply. The findings highlight the importance of solver selection in enhancing the management, scalability, and reliability of VPP optimization strategies, offering insights into the optimal combination of programming interfaces and solvers for efficient VPP operation.

I. INTRODUCTION

AROUND the world, Renewable Energy Sources (RESs) have taken advantage of the strong development of Distributed Energy Resources (DERs) [1]. A solution to this challenge is to aggregate the RESs, assuming that there may be unstable output and inconsistent generation of individual RES to emerge like a conventional generator with relatively stable output. Virtual Power Plants (VPPs) provide the potential solution for this problem, integrating Cyber-Physical Systems (CPS) to enhance the efficiency and coordination of various distributed energy resources [2].

The concept of VPPs represent a transformative approach in the energy sector, aiming to integrate various DERs such as renewable energy sources, controllable loads, and Energy Storage Systems (ESSs) into a cohesive and optimized network. This integration is facilitated through advanced software and hardware technologies, allowing for centralized control while

This article is a result of the Innovation Pact “NGS - New Generation Storage” (reference 58), co-financed by NextGeneration EU, through the Incentive System “Agendas para a Inovação Empresarial” (“Agendas for Business Innovation”), within the Recovery and Resilience Plan (PRR). The GECAD team authors acknowledge the work facilities and equipment provided by GECAD research center (UIDB/00760/2020) to the project team.

maintaining the autonomy of individual resources. The concept of VPPs is becoming increasingly important as the global shift towards RESs intensifies, requiring innovative solutions to manage the variability and decentralization inherent in these sources [3].

In this sense, based on strategies and technologies for monitoring, controlling and programming DERs, it is possible for a VPP to generate benefits such as decreasing the customer’s energy cost, reducing emissions, increasing energy efficiency, and asset control/optimization [4]. In this paper, a benchmark study for a VPP algorithm is addressed, where energy balance is guaranteed, the battery’s state of charge is kept within safe operating limits, and electrical appliances are shifted accordingly to the user’s requirements, all with a view to minimize operating costs. The scenario that served as the basis for this work was based on the model formulated in [5]. This deals with energy consumption management in the residential sector, as it is crucial to mitigate peak demand. It is based on reprogramming household appliances to change their load during peak hours, which significantly helps the grid. By combining the capabilities of ESS and load-shifting appliances, the Home Energy Management System (HEMS) can intelligently program and coordinate the operation of these devices to maximize the use of renewable energy resources and minimize electricity costs.

Moreover, this paper explores the development of a MILP algorithm designed to optimize the energy exchange among different houses within a VPP, focusing on the decision-making process regarding energy trading in specific time frames. The need for such optimization arises from the complex dynamics of energy demand and supply, the integration of RESs, and the economic considerations of buying and selling energy in the competitive market. The study delves into the performance of various commercial and open-source solvers in handling the optimization model. This comparative analysis is crucial for identifying the most effective computational tools for VPP optimization, taking into account factors such as computational efficiency, scalability, and the ability to handle VPP operations.

This paper is organized as follows: Section II presents the

literature review on VPPs, highlighting their significance, operational challenges, and the role of optimization algorithms. Section III describes the optimization algorithm used in VPP, detailing the mathematical model, objectives, parameters, and key constraints. Section IV discusses the practical implementation of the optimization service, including the architecture, data workflow, and comparison between commercial and open-source solvers. Section V presents the results of the study, focusing on the input/output data, test scenarios, implementation, and discussion. Finally, Section VI concludes the paper and suggests directions for future work.

II. LITERATURE REVIEW

The emergence of VPPs marks a significant evolution in the power system architecture, aggregating DERs such as distributed generation, energy storage, and flexible charging capabilities to provide vital grid services [6], [7]. The integration of VPPs into the energy system enhances efficiency, reliability, and sustainability, addressing the challenges posed by the increasing penetration of variable RESs like solar and wind energy. Since the early 2000s, the concept of VPPs has gained prominence as a solution to the variability and unpredictability of RESs, which, despite their clean and renewable nature, introduce issues related to system stability and reliability [2].

To further enrich our understanding of VPPs and their pivotal role in integrating DERs, it's crucial to explore the operational challenges they face in greater depth, particularly in terms of reliability, scalability, and security. VPPs, being at the forefront of the transition to a decentralized energy system, face unique challenges related to the intermittent nature of RESs, the complexity of managing diverse energy assets, and ensuring cybersecurity in an increasingly digitalized infrastructure. VPPs offer a strategic response to these challenges by enhancing system flexibility. They play a crucial role in balancing the generation from RESs with demand, providing essential grid services such as frequency re-balancing, load management, and participating in energy markets. This ability to mitigate the impacts of variable energy sources on the grid underlines the importance of developing advanced optimization algorithms for VPP operation [4].

When exploring the development and optimization of VPPs, the importance of Cyber-Physical Systems (CPS) in smart grids becomes evident [8]. The integration of such systems is crucial for the efficiency of VPPs, as it facilitates communication and coordination between the various DERs, from generation to consumption. This context lays the foundation for understanding VPPs not just as technological entities, but as complex energy ecosystems that require advanced optimization to operate effectively.

The optimization of VPPs, especially with regard to the control of microgrids, is an area of growing interest in the scientific community [9]. Effective control of microgrids, essential components of VPPs, plays a significant role in the management of DERs, highlighting the need to develop robust algorithms that can deal with the complexity and dynamism of

the energy system. The approach of using MILP to optimize DERs in microgrids illustrates the applicability of this methodology in solving complex optimization problems in VPPs [10]. The ability to model sizing, allocation and operation decisions in an integrated way offers a powerful tool to optimize the performance of VPPs, ensuring energy efficiency and sustainability. The objective is to minimize energy usage costs in a day-ahead operation. Furthermore, on the application of MILP in integrated power systems provides valuable insights into the different interfaces and solvers used in optimizing VPPs ([11]). This analysis highlights the diversity of available tools and the importance of selecting the most appropriate approach for each specific scenario, emphasizing the need for detailed benchmarks that evaluate the performance of these different solvers.

More recently, some works provide insights into the current trends and methodologies in VPP optimization, including the application of MILP models, consideration of trade markets, and the optimization of solvers to enhance VPP operations [12], [13]. Recent studies have underscored the importance of addressing these challenges through advanced optimization strategies, real-time data analytics, and robust cybersecurity protocols. For instance, the integration of machine learning algorithms for predictive analysis can significantly enhance the forecasting accuracy of renewable energy production and consumption patterns, thus improving the VPP's ability to balance supply and demand effectively [14], [15].

To ensure the resilience of VPP operations, research has also focused on developing sophisticated strategies that can dynamically adapt to changes in the energy market and regulatory environments. This includes the application of adaptive optimization algorithms that can accommodate multiple objectives, such as minimizing costs, maximizing the use of renewable energy, and ensuring grid stability. The contribution of this article lies in the application of a MILP algorithm developed specifically to optimize energy exchanges within VPP, addressing the complexity inherent in energy purchase and sale decisions taking into account the shifts of home appliances' energy management. By comparing the performance of various solvers and interfaces, this research not only fills an identified gap in the existing literature, but also offers practical guidelines for effectively implementing optimization solutions in VPPs. This approach allows for a deeper understanding of the capabilities and limitations of available solvers, guiding future research and development in optimizing VPPs to improve the sustainability and efficiency of the energy system.

III. OPTIMIZATION ALGORITHM IN VPP

The optimization algorithm for VPPs needs a comprehensive approach that not only considers the economic objectives, such as cost minimization and revenue maximization but also integrates technical constraints including battery storage capacities, renewable energy forecasts, and load demand variations.

A critical component of the VPP algorithm involves the formulation of a robust decision-making framework that can efficiently manage the scheduling of appliance usage, the charging and discharging of ESSs, the dispatch of distributed generation units, and the real-time bidding in electricity markets. This requires the inclusion of predictive models that use historical data and real-time inputs to forecast prices, generation capacity, and demand, thereby enabling more accurate and dynamic optimization. The description of the general model proposed in this work aims to support the management of the VPP. The notation used for the model is presented below, together with the mathematical formulation of the objectives, parameters, assumptions, and key constraints of the model.

The VPP operation is modeled using a MILP mathematical programming model that considers integer and continuous variables and all functions, objective and constraints, are linear. In terms of notation for indices and sets, different time periods ($t \in N_t$), different houses ($i \in N_i$) and different appliances ($l \in N_l$) were considered. N_t , N_i and N_l refer to the total number of periods, houses, and appliances respectively. It should be noted that the model in [5] refers to one household, however, as the aim is to study the optimization of a VPP, tests were carried out with several households.

This section presents the mathematical formulation used to model the problem of VPP energy management, which considers the scheduling of shift loads.

Equation 1 represents the objective function of the VPP, which is the calculation of energy costs:

$$EC = \sum_{i=1}^{N_i} \sum_{t=1}^{N_t} \left(P_{t,i}^{\text{buy}} \times \text{tou}_{t,i}^{\text{buy}} - P_{t,i}^{\text{sell}} \times \text{fit}_{t,i}^{\text{sell}} \right) \times \Delta t + c_i^{\text{fix}} \quad (1)$$

where EC represents the energy costs, $P_{t,i}^{\text{buy}}$ is the power purchased from the grid, $\text{tou}_{t,i}^{\text{buy}}$ is the time-of-use tariff for buying, $P_{t,i}^{\text{sell}}$ is the power sold to the grid, $\text{fit}_{t,i}^{\text{sell}}$ is the tariff for selling power, Δt represents the hourly adjust value, and c_i^{fix} is the fixed costs. Equation 2 represents the VPP energy balance:

$$P_{t,i}^{\text{buy}} + P_{t,i}^{\text{gen}} + P_{t,i}^{\text{dch}} = P_{t,i}^{\text{sell}} + \sum_{l=1}^{N_l} P_{t,i,l}^{\text{shift}} + P_{t,i}^{\text{load}} - P_{t,i}^{\text{ch}}, \quad \forall t \in N_t, \forall i \in N_i \quad (2)$$

where $P_{t,i}^{\text{gen}}$ is the power generated by the PV system, $P_{t,i}^{\text{dch}}$ is the power discharged from the battery, $P_{t,i,l}^{\text{shift}}$ is the power of shifted controllable loads, $P_{t,i}^{\text{load}}$ is the power consumed by the uncontrollable loads and $P_{t,i}^{\text{ch}}$ is the power charged to the battery. Equation 3-5 are used to simulate the buying and selling of electricity:

$$0 \leq P_{t,i}^{\text{buy}} \leq \overline{P_{t,i}^{\text{buy}}} \times Y_{t,i}^{\text{buy}}, \quad \forall t \in N_t, \forall i \in N_i \quad (3)$$

$$0 \leq P_{t,i}^{\text{sell}} \leq \overline{P_{t,i}^{\text{sell}}} \times Y_{t,i}^{\text{sell}}, \quad \forall t \in N_t, \forall i \in N_i \quad (4)$$

$$Y_{t,i}^{\text{buy}} + Y_{t,i}^{\text{sell}} \leq 1, \quad \forall t \in N_t, \forall i \in N_i \quad (5)$$

where $\overline{P_{t,i}^{\text{buy}}}$ and $\overline{P_{t,i}^{\text{sell}}}$ are the maximum limits for buying and selling electricity, and $Y_{t,i}^{\text{buy}}$, $Y_{t,i}^{\text{sell}}$ are binary variables indicating whether buying or selling actions are taken. In this case, the use of binary variables is important to ensure that only one action is performed in a given period. Equations 6-9 are used to emulate the battery behavior:

$$0 \leq P_{t,i}^{\text{ch}} \leq \overline{P_{t,i}^{\text{ch}}} \times Y_{t,i}^{\text{ch}}, \quad \forall t \in N_t, \forall i \in N_i \quad (6)$$

$$0 \leq P_{t,i}^{\text{dch}} \leq \overline{P_{t,i}^{\text{dch}}} \times Y_{t,i}^{\text{dch}}, \quad \forall t \in N_t, \forall i \in N_i \quad (7)$$

$$Y_{t,i}^{\text{ch}} + Y_{t,i}^{\text{dch}} \leq 1, \quad \forall t \in N_t, \forall i \in N_i \quad (8)$$

$$\underline{SoC_{t,i}^{\text{bat}}} \leq SoC_{t,i}^{\text{bat}} \leq \overline{SoC_{t,i}^{\text{bat}}}, \quad \forall t \in N_t, \forall i \in N_i \quad (9)$$

where, $\overline{P_{t,i}^{\text{ch}}}$ represents the maximum limit for battery charge, $Y_{t,i}^{\text{ch}}$ is a binary variable associated to the battery charge action, $\overline{P_{t,i}^{\text{dch}}}$ denotes the maximum value for battery discharge, $Y_{t,i}^{\text{dch}}$ indicates the binary variable associated to the discharge action, $SoC_{t,i}^{\text{bat}}$ represents the state of charge of the battery, $\underline{SoC_{t,i}^{\text{bat}}}$ and $\overline{SoC_{t,i}^{\text{bat}}}$ are the minimum and maximum limit for the state of charge of the battery. Equation 10 is used to calculate the balance of the battery for period $t=1$ and equation 11 to calculate the balance in the remaining periods:

$$SoC_{t,i}^{\text{bat}} = SoC_i^{\text{bat init}} + (P_{t,i}^{\text{ch}} - P_{t,i}^{\text{dch}})\Delta t, \quad t = 1, \forall i \in N_i \quad (10)$$

$$SoC_{t,i}^{\text{bat}} = SoC_{t-1,i}^{\text{bat}} + (P_{t,i}^{\text{ch}} - P_{t,i}^{\text{dch}})\Delta t, \quad \forall t \in [2, N_t], \forall i \in N_i \quad (11)$$

where $SoC_i^{\text{bat init}}$ is the state of charge value for the first instant. Equation 12 presents the shift power loads calculation:

$$P_{t,i,l}^{\text{shift}} = P_{t,i,l}^{\text{controllable load}} \cdot z_{t,i,l}, \quad \forall t \in N_t, \forall i \in N_i, \forall l \in N_l \quad (12)$$

where, $P_{t,i,l}^{\text{controllable load}}$ represents the power for the controllable load and $z_{t,i,l}$ indicates whether the load is turned on or off. Equations 13 and 14 are used to control and shift the controllable loads:

$$z_{t,i,l} \leq t_{t,i,l}^{\text{on}}, \quad t = 1, \forall i \in N_i, \forall l \in N_l \quad (13)$$

$$z_{t,i,l} \leq z_{t-1,i,l} + t_{t,i,l}^{\text{on}}, \quad \forall t \in [2, N_t], \forall i \in N_i, \forall l \in N_l \quad (14)$$

where, $t_{t,i,l}^{\text{on}}$ represents a binary variable that indicates the moment when the load is turned on. Equation 15 ensures

that shift loads will be activated during the predefined period number, while equation 16 ensures that each shift load will only be activated once:

$$\sum_{t=1}^{N_t} z_{t,i,l} = P_{t,i,l}^{\text{controllable load periods}}, \quad \forall i \in N_i, \forall l \in N_l \quad (15)$$

$$\sum_{t=1}^{N_t} t_{t,i,l}^{\text{on}} = 1, \quad \forall i \in N_i, \forall l \in N_l \quad (16)$$

where $p_{t,i,l}^{\text{controllable load periods}}$ represents the total number of periods that each shift load needs to be activated. Equation 17 presents the limits for binary variables.

$$\{Y_{t,i}^{\text{buy}}, Y_{t,i}^{\text{sell}}, Y_{t,i}^{\text{ch}}, Y_{t,i}^{\text{dch}}, z_{t,i,l}, t_{t,i,l}^{\text{on}}\} \in \{0, 1\} \quad (17)$$

IV. PRACTICAL IMPLEMENTATION

The optimization service plays a pivotal role in the operational framework of the system. It is designed as a modular service, which intakes operational data from both external sources, processes this data through a mathematical optimization model, and receive the outputs and directives for the operational control. The core of this service is the optimization algorithm, which relies on solver software to find optimal solutions within a predefined time constraint. Fig. 1 presents the architecture and data workflow of the optimization service.

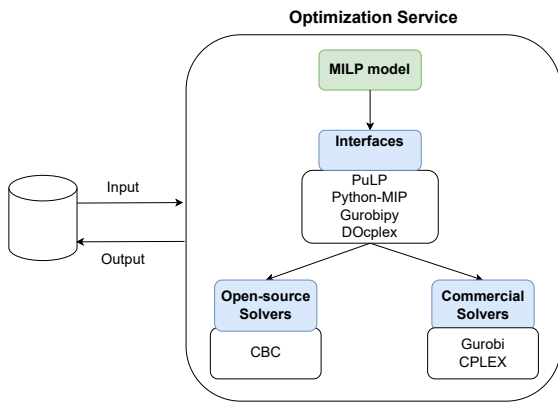


Fig. 1. Architecture and data flow for the optimization service.

The service fetches data through an abstract class, meaning any database system could be integrated if a concrete implementation of the abstract class is provided. This class defines an interface with a single function, which takes the input for the algorithm (defined in Table I) and returns a solution as an output (defined in Table II). The optimization service can be considered as a single worker; i.e., at any given moment, it can only process one single input dataset using one single solver. This does not imply that no parallelization is achievable, however, as it is up to the solvers to adopt such parallel computing strategies internally. If the optimization of multiple input datasets is required, then multiple instances of

the optimization service can be deployed, every single one of them with a different solver, if desired. It should be noted that there is no automatic choice of solver; the solver must be chosen manually during the configuration of the service.

The practical implementation of VPP optimization services requires a critical evaluation of the available solver technologies, both commercial and open-source. Our comparison focuses on aspects such as computational efficiency, scalability, ease of integration, and cost-effectiveness. In the realm of optimization problems, especially those as complex and dynamic as those encountered in VPP operations, the choice of solver can significantly impact the efficiency and reliability of the solution process. While open-source solvers offer accessibility and flexibility, commercial solvers typically provide superior performance in terms of computational speed and problem-solving capabilities. Two of the most renowned commercial solvers are Gurobi and CPLEX [16], [17], both of which have established a reputation for their robustness in handling large-scale optimization problems across various industries, including energy management and optimization in VPPs.

To achieve optimal performance, it is important to examine the solver behavior through the log file analysis. If feasible, it is advisable to explore various solvers, in order to expand the range of options. However, it is crucial to exercise caution when relying solely on recommendations for well-established, state-of-the-art solvers, particularly without conducting practical computational experiments. Additionally, it is essential to take into account the complexity of the modeling environment/language and estimate the amount of time required to complete the modeling phase.

A. Commercial Solvers

Commercial solvers play a pivotal role in the efficient handling of complex optimization problems. Below are two solvers, Gurobi and CPLEX, renowned for their advanced algorithms and high-performance capabilities.

- **Gurobi:** Developed by Gurobi Optimization, LLC, the Gurobi Optimizer is a state-of-the-art solver for a wide range of optimization problems, including Linear Programming (LP), MILP, Quadratic Programming, and Mixed-integer Quadratic Programming. Gurobi is celebrated for its high-performance computing capabilities, scalability, and comprehensive support for programming languages and development environments. It integrates advanced algorithms that can be tuned for specific problem types, ensuring optimal performance and accuracy in solving complex optimization models [16].
- **CPLEX:** IBM ILOG CPLEX Optimization Studio encompasses solvers for LP, MILP, QP, and MIQP problems. CPLEX Optimizer stands out for its powerful pre-solve and cutting-plane algorithms, which efficiently reduce problem size and complexity, significantly speeding up the solution process. Its parallel processing capabilities allow it to leverage multiple processors to handle intricate problems more efficiently, making it particularly suitable

for the demanding optimization tasks associated with managing VPP operations. CPLEX also offers a flexible API and is supported by a wide range of programming languages, facilitating its integration into custom applications [17].

B. Open-Source Solvers

CBC (COIN-OR Branch-and-Cut) is highlighted as a notable open-source solver for addressing various optimization problems [18]. CBC has many key advantages, particularly its cost-effectiveness due to the absence of licensing fees. This attribute is especially valuable for enabling the application of advanced optimization techniques without incurring the high costs associated with commercial software. Such economic efficiency promotes broader adoption and experimentation with optimization models, fostering innovation and research in energy management and beyond. It supports various modeling languages and interfaces, such as PuLP and Python-MIP, facilitating its incorporation into existing systems and workflows. This ease of integration significantly reduces development time and complexity, rendering CBC a practical choice for developers and researchers working on optimization problems. Its compatibility with modeling languages and interfaces ensures that CBC can seamlessly fit into diverse computational environments, enabling the formulation and solving of MILP models with efficiency and precision. Therefore, the following abstraction interfaces were chosen:

- **PuLP:** It is a LP open-source library written in Python. It serves as a modeling layer for LP and MILP problems. The choice of PuLP is motivated by its simplicity in defining decision variables, objectives, and constraints directly in Python, and its compatibility with multiple solvers (CBC, GLPK, CPLEX, and Gurobi), offering flexibility in solver selection based on the problem at hand and available licenses. This makes PuLP ideal for scenarios requiring straightforward modeling and diverse solver option [19].
- **Python-MIP:** Focused on MILP problems, Python-MIP is particularly noted for its performance and advanced features, essential for solving large and complex MILP efficiently. The choice of Python-MIP was driven by the need for a tool that provides deeper integration with MILP solver technologies, enabling more sophisticated problem-solving strategies. This tool provides specialized features for MILP that are not as readily accessible in the other tools [20].

V. RESULTS

This section presents the results from the benchmark study of solvers for VPP optimization, detailing the performance and efficiency of different solver technologies in managing and optimizing VPP operations.

A. Input/Output Data

As detailed in Section III, the model takes in consideration several constraints and decision variables, where only some are

provided as input and the rest as output from the solver. The input data used in this study were collected from real historical data of residential energy consumption and local photovoltaic generation. This information was obtained through continuous measurements in households utilizing HEMS. These data include appliance load profiles, photovoltaic energy production, and electricity tariffs [5]. The input data, as outlined in the Table I, includes a variety of parameters essential for the model to perform accurately.

These parameters encompass the number of periods (Nt), the number of houses (Ni), time interval between periods (Δt), and various other technical and economic factors that influence the VPP energy management decisions. The range of values for these parameters is meticulously defined to accommodate the diversity in VPP configurations and operational scenarios.

The output data, presented in Table II, details the results produced by the optimization model, which are critical for making informed decisions regarding energy transactions within the VPP.

These outputs are crucial in strategizing the VPP operations to enhance efficiency, reduce costs, and maintain energy balance.

This distinction between input and output data underscores the comprehensive approach adopted in the optimization model, where various economic and technical aspects are taken into consideration to optimize the VPP operations effectively. The model ability to process a wide range of input parameters and produce actionable outputs facilitates the effective management of energy resources, thereby contributing to the overall efficiency and sustainability of the VPP ecosystem.

B. Testing Environment

All tests were executed on a machine with the following hardware/software specifications:

- **CPU:** Intel Xeon E5-2686 v4 (only 4 vCPU)
- **RAM:** 16 GB
- **OS:** Ubuntu 24.04 LTS (virtualized), running Linux 6.8.0

As for Python and optimization libraries, the following versions were used:

- CPython 3.10.14
- PyPy 7.3.16 (implements the Python 3.10.14 standard)
- Gurobi 11.0.2
- CPLEX 22.1.1.0 (with DOcplex 2.27.239)
- Python-MIP 1.15.0
- PuLP 2.7.0

To measure the scalability of the algorithm, four different test scenarios were considered: 1 house, micro VPP (2 houses), small VPP (4 houses) and community VPP (8 houses). A fixed number of 96 periods with 15 minutes each was set, to allow a full day-ahead optimization. Furthermore, it is important to note that each house has controllable appliances for energy management, more specifically 12 appliances (including, for example, clothes washing machines, hair dryers, coffee makers, and phone chargers, among others). Each

TABLE I
INPUT DATA OF THE ALGORITHM.

Parameter	Designation	Value Range	Unit
N_t	Number of periods	96	-
N_i	Number of houses	2 – 8	-
N_l	Number of controllable loads	12	-
Δt	Time interval between periods	0.25	h
$to_{t,i}^{buy}$	Price for buying electricity	0.1034 – 0.2314	€/kWh
$fi_{t,i}^{sell}$	Price for selling electricity	0.045	€/kWh
c_i^{fix}	Fixed costs	0.2197 – 0.6249	€
$P_{t,i}^{buy}$	Maximum buy amount	4.6 – 13.8	kWh
$P_{t,i}^{sell}$	Maximum sell amount	4.6 – 13.8	kWh
$P_{t,i}^{ch}$	Maximum charge amount	0 – 5	kWh
$P_{t,i}^{dch}$	Maximum discharge amount	0 – 5	kWh
$P_{t,i}^{gen}$	Generated energy forecast	0 – 8.474	kWh
$P_{t,i}^{load}$	Consumed energy forecast	0.052 – 9.822	kWh
$P_{t,i,l}^{controllable\ load}$	Controllable load power	0.01 – 5.20	kWh
$P_{t,i,l}^{controllable\ load\ periods}$	Number of periods that each shift load needs to be activated	1 – 8	-
$SoC_i^{bat\ init}$	Initial state of charge of the battery	0 – 1.92	kWh
$SoC_{t,i}^{bat}$	Minimum SoC of the battery	0 – 1.824	kWh
$SoC_{t,i}^{bat}$	Maximum SoC of the battery	0 – 9.6	kWh

TABLE II
OUTPUT DATA OF THE ALGORITHM.

Parameter	Designation	Unit
$P_{t,i}^{buy}$	Energy amount to be bought	kWh
$P_{t,i}^{sell}$	Energy amount to be sold	kWh
$P_{t,i}^{ch}$	Energy amount to be charged	kWh
$P_{t,i}^{dch}$	Energy amount to be discharged	kWh
$SoC_{t,i}^{bat}$	Resulting state of charge of the battery	kWh
$P_{t,i,l}^{shift}$	Power of shifted loads	kWh
$Y_{t,i}^{buy}$	Indicates whether energy was bought or not	Binary
$Y_{t,i}^{sell}$	Indicates whether energy was sold or not	Binary
$Y_{t,i}^{ch}$	Indicates whether energy was charged or not	Binary
$Y_{t,i}^{dch}$	Indicates whether energy was discharged or not	Binary
$z_{t,i,l}$	Indicates whether the load is on or not	Binary
$t_{t,i,l}^{on}$	The period when the load is first turned on	Binary

house has the same number of appliances, and the idea is to implement strategies to monitor and schedule appliance usage in a manner that reduces overall energy consumption and minimizes electricity costs, regardless of the appliances themselves.

Each optimization library was tested on each one of the four test scenarios, only accounting for the time spent using

the solver (i.e., preparing the problem and optimizing); the fetching of the data and subsequent output were not considered. The data was collected through the measurement of five separate executions on sequentially-running processes.

This study does not aim to provide a direct comparison between PyPy—an alternative Python implementation with a just-in-time (JIT) compiler—and CPython; the two runtimes are merely presented with the aim of analyzing how each optimization library performs in their respective environments. As noted by the PyPy developers, the usage of short-lived processes and external libraries—which is, in essence, the workload featured in this article—is not well-suited for JIT compilation, and, as such, performance benefits will not be visible¹. In addition, since PyPy has no support for the official Gurobi and DCOplex libraries, those tests were omitted from the final results.

C. Implementation and Discussion

Implementing an optimization service for VPP requires careful consideration of several factors beyond the selection of a solver. These include the complexity of the optimization model, the scalability of the solution (to accommodate varying sizes of VPP networks), and the integration of the optimization service with existing data sources and control systems. The practical experiences of deploying commercial and open-source solvers in VPP optimization tasks have shown that,

¹<https://www.pypy.org/features.html#speed> (accessed 24 May 2024)

TABLE III
EXECUTION TIME (IN SECONDS) OF THE SOLVERS, IN TERMS OF HOUSE COUNT.

Runtime	Solver	House (1)	Micro VPP (2)	Small VPP (4)	Community VPP (8)	
CPython	docplex	0.313	1.483	2.543	9.858	
	gurobi	0.348	1.314	2.375	4.541	
	pulp:cbc	1.174	41.402	118.666	7272.551	
	pulp:cplex	0.540	1.868	3.988	20.937	
	pulp:gurobi	0.406	1.382	2.674	5.670	
	python-mip:cbc	0.798	24.499	110.241	2386.632	
	python-mip:gurobi	0.301	1.258	2.058	4.409	
PyPy	pulp:cbc	1.526	41.842	119.127	7279.646	
	pulp:cplex	1.002	2.444	4.493	20.965	
	pulp:gurobi	0.872	1.867	3.044	5.512	
	python-mip:cbc	1.244	24.979	110.855	2381.551	
		python-mip:gurobi	0.716	1.693	2.419	4.391

while both types of solvers deliver different performance, their suitability can vary depending on the specific requirements of the task at hand, such as the problem size, the complexity of constraints, and the computational resources available. For example, the computational experiments revealed a significant performance difference between commercial and open-source solvers for MILP problems, even in small examples. This can be attributed to algorithmic optimization, parallel processing, advanced presolve techniques and cutting-edge features. The results can be seen in Table III. The table is split into two sections for CPython and PyPy runtimes.

Observing the previous table, it is also possible to see that certain parameters influence the exploratory tests of different interfaces and solvers:

- **Solver performance variability:** The discussion table highlights significant variability in solver performance, with Gurobi and CPLEX outperforming CBC in terms of speed and reliability, especially as the problem size increases. This aligns with the expectations set in the literature review, where commercial solvers are often recognized for their superior optimization capabilities and efficiency in handling large-scale, complex problems.
- **Impact of problem size on solver efficiency:** As the number of houses increases, the computational complexity of the optimization problem escalates, challenging the solvers capabilities to find optimal solutions within reasonable time frames. The table shows a clear trend of increasing execution time with the number of houses, highlighting the importance of choosing a solver that scales well with problem size for practical VPP applications.
- **Optimization under real-world constraints:** The results underscore the necessity of employing solvers that can effectively handle the intricate constraints typical of VPP operations, such as electricity costs, energy balance and battery management and the management of twelve individual appliances through shift actions. The ability of a solver to navigate these constraints efficiently is critical

TABLE IV
DETAILED EXECUTION TIME (IN SECONDS) OF THE SOLVERS FOR ONE HOUSE.

Runtime	Solver	Mean ± stdev	Min	Max
CPython	docplex	0.313 ± 0.001	0.312	0.314
	gurobi	0.348 ± 0.007	0.339	0.358
	pulp:cbc	1.174 ± 0.007	1.166	1.185
	pulp:cplex	0.540 ± 0.008	0.533	0.552
	pulp:gurobi	0.406 ± 0.003	0.402	0.409
	python-mip:cbc	0.798 ± 0.007	0.791	0.810
	python-mip:gurobi	0.301 ± 0.003	0.297	0.305
PyPy	pulp:cbc	1.526 ± 0.013	1.517	1.548
	pulp:cplex	1.002 ± 0.020	0.985	1.031
	pulp:gurobi	0.872 ± 0.088	0.825	1.029
	python-mip:cbc	1.244 ± 0.008	1.236	1.254
		python-mip:gurobi	0.716 ± 0.006	0.707

for optimizing VPP performance, minimizing operational costs, and ensuring energy supply meets demand. Here all interfaces showed the ability to deal with the problem in question.

The results in the table demonstrate that Python-MIP, combined with the Gurobi solver, achieves the best results. This outcome highlights several critical factors in the context of optimization: solver compatibility and efficiency, algorithmic enhancements and leveraging commercial solver strengths.

In order to have a clear understanding of the results, the variation in the performance of the solvers is represented below for an example case of one house. In this sense, Table IV provides a detailed analysis of the execution times for different solvers in a VPP optimization scenario with just one house. Each solver performance is measured in terms of the mean execution time with its standard deviation, as well as the minimum and maximum execution times observed.

In the CPython runtime section, the “python-mip:gurobi” solver exhibits the best average performance with a mean execution time of 0.301 seconds. The “docplex” solver follows closely, indicating that commercial solvers have the upper hand in performance. CBC, both under PuLP and Python-MIP, offers the worst level of performance of the list, with Python-MIP beating PuLP by a slight margin.

Under the PyPy runtime, all solvers follow the same ranking in regard to execution times. Particularly, “python-mip:gurobi” achieves the best performance with a mean execution time of 0.716 seconds.

In order to visually capture the results from Table IV, a box plot visualization (Fig. 2) was generated.

The box plot provides is a graphical tool to represent the variation in observed statistical data using quartiles (e.g., minimum, maximum, and median). In turn, outliers can be plotted as individual points. This visual evidence supports an analysis in terms of consistency and variability in solver performance, providing empirical data on their performance in a VPP context.

In terms of the main observations made, it is worth highlighting that:




Fig. 2. Box plot with the elapsed time for optimization by various solvers for a scenario with one house setup.

- commercial solvers offer the best level of performance, with low execution times even in more demanding workloads;
- CBC is not suitable for real-time usage, as the Community VPP scenario takes between 30 minutes and 2 hours to complete;
- solvers executed under CPython generally have a lower average solving time compared to those run under PyPy, as the test workload is not suitable for JIT compilation;
- the spread of the data for each solver is quite contained.

This visualization effectively highlights the efficiency and consistency differences between solver implementations and the impact of the Python runtime environment on their performance. These results clearly illustrate that the choice of both solver and runtime can significantly impact performance in VPP optimization tasks. It also indicates that, while open-source solvers like CBC are slower, they may offer a cost-effective alternative to commercial solvers, especially when used for smaller VPP sizes.

In essence, commercial solvers show superior performance, suggesting their suitability for larger and more complex VPP networks. However, the choice of solver also needs to consider factors such as licensing costs, software compatibility, and the availability of technical support. While commercial solvers like Gurobi and CPLEX offer high performance and robust support, their cost might not be justifiable for all project scales. Conversely, open-source solvers like CBC provide a viable alternative with flexibility and customization options, although with potential trade-offs in terms of execution speed and solution optimality. In this sense, it becomes important to discuss the advantages of using open-source solvers. The most evident advantage of using CBC, or any open-source solver, is the absence of licensing fees. This cost efficiency can be a critical factor in enabling the use of advanced optimization techniques without the financial burden associated with commercial software. This democratization of access allows

for a wider adoption and experimentation with optimization models, fostering innovation and research in the field of energy management and beyond. Furthermore, open-source software offers unparalleled flexibility and customization opportunities. Users can modify the source code to tailor the solver to their specific needs, optimize its performance for particular types of problems, or even contribute improvements back to the community. This adaptability is particularly beneficial in the rapidly evolving domain of VPPs, where unique and complex optimization challenges can necessitate specialized solver functionalities. In addition, open-source projects benefit from the support of a broad and active community. Users and developers can collaborate, share knowledge, and offer support through forums, repositories, and direct contributions. This collective intelligence can accelerate problem-solving and innovation. Additionally, the transparency of open-source software ensures that the algorithms and methodologies employed are fully visible and open to scrutiny, fostering trust and understanding among users. Finally, the CBC solver, being a part of the Computational Infrastructure for Operations Research (COIN-OR) project, is designed with compatibility and integration in mind. It can be used with various modeling languages and interfaces, such as PuLP and Python-MIP, facilitating its incorporation into existing systems and workflows. This ease of integration can significantly reduce development time and complexity, making it a practical choice for a wide range of optimization problems. It should be emphasized that the purpose was not to observe the effects of the optimization results on the VPPs, but rather to theoretically evaluate the performance of the interfaces.

Thus, for many practical applications, CBC and other open-source solvers can offer sufficient performance and capabilities, especially when the problem is well-structured and falls within the solver optimization strengths. For VPP operators and developers, this suggests a strategic approach to solver and interface selection, taking into consideration not only the mathematical and computational capabilities but also the execution environment. The findings encourage further exploration and benchmarking of different combinations to identify the most effective setup for specific VPP optimization scenarios.

VI. CONCLUSIONS AND FUTURE WORK

The research presented in this document underscores the critical role that solver selection plays in the effective management and optimization of VPPs. Through a detailed comparison of solver performance across various scenarios. The study offers valuable insights that pave the way for the development of more efficient, scalable, and reliable VPP optimization strategies to optimize house energy management considering shifting of electric appliances.

The findings point to Python-MIP and Gurobi as a particularly promising combination for achieving high efficiency in VPP optimization, especially in the context of any sized networks. However, it is crucial for stakeholders to conduct a thorough assessment of a project specific requirements when choosing between commercial and open-source solvers.

Commercial solvers like Gurobi and CPLEX, known for their robustness, appear particularly well-suited for larger, more complex VPP systems, while open-source solvers like CBC are highlighted for their cost efficiency and flexibility, making them ideal for scenarios with small to medium-sized networks where these factors are prioritized. The performance, scalability, and support services of the solver are critical factors that influence the efficiency and reliability of the VPP optimization service.

Despite the strengths of the study, including the comprehensive evaluation of solver performance, there are limitations to consider. Future research should expand on these findings, exploring the scalability of these solutions for larger VPP networks and integrating machine learning techniques to better predict energy demand and production, thereby enhancing the operational efficiency of VPPs. Furthermore, experiments may be carried out that include JIT performance tests as future work. Moreover, meta-heuristic techniques and high-performance computing could be works that leverage the strengths of multiple approaches might provide innovative solutions for VPP optimization in a wider range of scenarios.

REFERENCES

- [1] K. O. Adu-Kankam and L. M. Camarinha-Matos, "Renewable energy communities or ecosystems: An analysis of selected cases," *Heliyon*, vol. 8, no. 12, 2022. doi: <https://doi.org/10.1016/j.heliyon.2022.e12617>
- [2] N. Naval and J. M. Yusta, "Virtual power plant models and electricity markets - a review," *Renewable and Sustainable Energy Reviews*, vol. 149, p. 111393, 2021. doi: <https://doi.org/10.1016/j.rser.2021.111393>
- [3] X. Wang, Z. Liu, H. Zhang, Y. Zhao, J. Shi, and H. Ding, "A review on virtual power plant concept, application and challenges," in *2019 IEEE Innovative Smart Grid Technologies - Asia (ISGT Asia)*, 2019. doi: [10.1109/ISGT-Asia.2019.8881433](https://doi.org/10.1109/ISGT-Asia.2019.8881433) pp. 4328–4333.
- [4] H. M. Rouzbahani, H. Karimipour, and L. Lei, "A review on virtual power plant for energy management," *Sustainable energy technologies and assessments*, vol. 47, p. 101370, 2021. doi: <https://doi.org/10.1016/j.seta.2021.101370>
- [5] R. Faia, P. Faria, and Z. Vale, "Optimizing house energy management with milp considering shifting of electric appliances and battery storage system," in *2023 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, 2023. doi: [10.1109/ISGTEUROPE56780.2023.10407660](https://doi.org/10.1109/ISGTEUROPE56780.2023.10407660) pp. 1–5.
- [6] E. A. Bhuiyan, M. Z. Hossain, S. Mueyeen, S. R. Fahim, S. K. Sarker, and S. K. Das, "Towards next generation virtual power plant: Technology review and frameworks," *Renewable and Sustainable Energy Reviews*, vol. 150, p. 111358, 2021. doi: <https://doi.org/10.1016/j.rser.2021.111358>
- [7] S. M. Nosratabadi, R.-A. Hooshmand, and E. Gholipour, "A comprehensive review on microgrid and virtual power plant concepts employed for distributed energy resources scheduling in power systems," *Renewable and Sustainable Energy Reviews*, vol. 67, pp. 341–363, 2017. doi: <https://doi.org/10.1016/j.rser.2016.09.025>
- [8] D. Ilic, P. G. Da Silva, S. Karnouskos, and M. Griesemer, "An energy market for trading electricity in smart grid neighbourhoods," in *2012 6th IEEE international conference on digital ecosystems and technologies (DEST)*. IEEE, 2012. doi: [10.1109/DEST.2012.6227918](https://doi.org/10.1109/DEST.2012.6227918) pp. 1–6.
- [9] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. A. Cañizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke *et al.*, "Trends in microgrid control," *IEEE Transactions on smart grid*, vol. 5, no. 4, pp. 1905–1919, 2014. doi: [10.1109/TSG.2013.2295514](https://doi.org/10.1109/TSG.2013.2295514)
- [10] S. Mashayekh, M. Stadler, G. Cardoso, and M. Heleno, "A mixed integer linear programming approach for optimal der portfolio, sizing, and placement in multi-energy microgrids," *Applied Energy*, vol. 187, pp. 154–168, 2017. doi: <https://doi.org/10.1016/j.apenergy.2016.11.020>
- [11] J. Mohamed, A. Muqbel, A. T. Al-Awami, and I. Elamin, "Optimal demand response bidding and pricing mechanism in distribution network: application for a virtual power plant," in *2018 IEEE Industry Applications Society Annual Meeting (IAS)*. IEEE, 2018. doi: [10.1109/IAS.2018.8544514](https://doi.org/10.1109/IAS.2018.8544514) pp. 1–8.
- [12] G. Weishang, W. Qiang, L. Haiying, and W. Jing, "A trading optimization model for virtual power plants in day-ahead power market considering uncertainties," *Frontiers in Energy Research*, vol. 11, p. 1152717, 2023. doi: <https://doi.org/10.3389/fenrg.2023.1152717>
- [13] Y. Gao, L. Gao, P. Zhang, and Q. Wang, "Two-stage optimization scheduling of virtual power plants considering a user-virtual power plant-equipment alliance game," *Sustainability*, vol. 15, no. 18, p. 13960, 2023. doi: [10.3390/su151813960](https://doi.org/10.3390/su151813960)
- [14] G. Ruan, D. Qiu, S. Sivaranjani, A. S. Awad, and G. Strbac, "Data-driven energy management of virtual power plants: A review," *Advances in Applied Energy*, p. 100170, 2024. doi: <https://doi.org/10.1016/j.adapen.2024.100170>
- [15] W. Nafkha-Tayari, S. Ben Elghali, E. Heydarian-Forushani, and M. Benbouzid, "Virtual power plants optimization issue: A comprehensive review on methods, solutions, and prospects," *Energies*, vol. 15, no. 10, p. 3607, 2022. doi: [10.3390/en15103607](https://doi.org/10.3390/en15103607)
- [16] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2021.
- [17] S. Nickel, C. Steinhardt, H. Schlenker, and W. Burkart, *Decision Optimization with IBM ILOG CPLEX Optimization Studio: A Hands-On Introduction to Modeling with the Optimization Programming Language (OPL)*. Springer Nature, 2022.
- [18] C. I. O. R. G. (COIN-OR), *CBC (Coin-or Branch and Cut) Solver Documentation*, 2011. [Online]. Available: <https://projects.coin-or.org/Cbc>
- [19] S. Mitchell *et al.*, *PuLP: A Linear Programming Toolkit for Python*, 2005. [Online]. Available: <https://coin-or.github.io/pulp/>
- [20] T. P.-M. Contributors, "Python-MIP: Modeling and solving mixed-integer linear programming problems," 2019, available at <https://www.python-mip.com/>.

Quality Control of Body Measurement Data Using Linear Regression Methods

Janis Bicevskis
0000-0001-5298-9859
University of Latvia
Raina boulevard 19,
Riga, LV-1586, Latvia
Email:
Janis.Bicevskis@lu.lv

Zane Bicevska
0000-0002-5252-7336
DIVI Grupa Ltd
Fridriha Candra str. 1,
Riga, LV-1046, Latvia
Email:
Zane.Bicevska@di.lv

Edgars Diebelis
0000-0002-5950-9915
DIVI Grupa Ltd
Fridriha Candra str. 1,
Riga, LV-1046, Latvia
Email:
Edgars.Diebelis@di.lv

Liva Purina
0009-0008-0339-5956
DIVI Grupa Ltd
Fridriha Candra str. 1,
Riga, LV-1046, Latvia
Email:
liva.purina@patterns.di.lv

Abstract—Body measurement data are inherently inaccurate and quite error-prone due to manual measurement and data collection. In this study, professionally collected and self-collected body measurement data were used to investigate to what extent potentially erroneous data can be identified during collection by utilizing the anthropologically given correlation of body measurements. The study specifically uses a dataset created within the framework of a project for made-to-measure pattern creation, consisting of data from 2053 female individuals with up to 52 recorded body measurements. Using linear regression, a method for validating the collected data is defined, wherein potentially inconsistent data are identified based on tolerance intervals. The tolerance intervals calculated within the study are specific to the particular application and the personal data used in the study. The outlined method is applicable to almost any set of manually collected body data in at least the triple-digit range, enabling the identification of probable data errors already during their collection.

Index Terms—linear regression, body measurement assessment, data quality in pattern generation.

I. INTRODUCTION

WITH the advancement of information technology capabilities, there are more and more specialized industry solutions that create, process, and use various types of graphic objects, such as photographs, drawn images, drawings, sketches, patterns, and so on. One type of these objects is technical drawings composed of polygons. The most typical examples are part drawings, geographic maps, and clothing patterns.

If polygon-based objects are created automatically (generated with the help of specialized programs and scripts), two types of problems arise: (A) How to ensure that the created object meets the established requirements (is correct)? (B) How to be confident that, even in the case of script changes,

the resulting output is as correct as in the previous version (continuity is maintained)?

A. Problem Identification

Translating the problem statement to the specific application intended for testing the study results — automated generation of clothing patterns — the described problem statement is reduced to a series of derivative questions: (a1) How to ensure that the generated pattern fits the clothing intended for a specific person?, (a2) How to ensure that the components of the generated pattern are compatible with each other, i.e., that the respective garment can be sewn together properly?, (b1) How to detect systematic deviations from the standard if the pattern generation script has been modified?, (b2) How to ensure that the pattern generation script is functional in all its branches (completes the work as desired and produces the necessary result)?

Before the digitalization era, problems of type (a) were solved by sewing and trying on clothing prototypes for people of specific sizes. As long as patterns are generated for fixed measurement sets (standard sizes), sewing prototypes, though resource-intensive, would still be technically feasible. However, in cases where individual patterns are created for a wide variety of measurement sets, this is no longer feasible even theoretically — the number of potential measurement combinations is virtually infinite.

One solution could be finding equivalence classes where the included sets of measurements behave similarly for a specific type of clothing. Another solution could involve elements of image recognition and automated comparison to "spot" extreme and boundary cases. In this study, the main focus is on assessing the consistency of input data, namely, the interrelations of body measurements. If a mechanism has been developed that reliably identifies atypical deviations (errors) in the data, in the next steps, these data could be used for machine learning and evaluating the accuracy of image recognition.

Problems of type (b) primarily relate to regression testing in the field of software engineering—specific quality criteria

This work has been conducted within the research project "Competence Centre of Information and Communication Technologies" of The Recovery and Resilience Facility, contract No. 5.1.1.2.i.0/1/22/A/CFLA/008 signed between IT Competence Centre and Central Finance and Contracting Agency, Research No. 1.7 "Technology for quality assessment of graphical objects described using polygons"

need to be identified that, on one hand, correspond to the task at hand and, on the other hand, are automatable and verifiable. The need for automated testing also arises from the very large number of combinations to be tested—for instance, to test the correctness of 160 scripts on 150 body profiles, at least 25,000 test cases would be required, which is practically impossible to execute manually.

Regression testing is a well-researched topic in scientific literature. However, specifically regarding quasi-continuous objects (such as graphic objects composed of polygons), there is a need for a critical evaluation of existing concepts. This aspect will not be considered in this study; however, it should be noted that one of the solutions for regression testing is given in [1]. The essence of the solution is as follows. Unlike traditional systems regression testing where test cases are constructed according to system specifications, in this case it is proposed to accumulate use cases from previous system applications. Each use case accepted by a customer, or an expert may become a regression test case, it ensures stability of the system as all previous use cases must work correctly after changes in the system. This will not only save the resources needed for preparation of test cases. It also ensures that the system is maintained in good quality as all previous use cases are run repeatedly. Customer-accepted test cases differ from the, often unrealistic, test cases created by testers according to system requirements specification.

B. Main Idea of Solution

To address the first research question (a1) - "How to ensure that the generated pattern fits the clothing intended for a specific person?" - it is necessary to assess the quality of human body measurements that will be used for pattern generation. With qualitative body measurements, we understand those that reflect specific measurements of a person's body part with an accuracy of up to ± 1 cm, and which, when used, can be used to automatically generate a pattern according to which clothing can be sewn to visually and functionally meet the requirements of the garment manufacturing industry — fitting snugly to the body, avoiding unwanted fabric gathers and protrusions, being symmetrical, allowing freedom of movement, and so on.

In the clothing manufacturing industry, automated individualized pattern creation for each client, without using a standard pattern base, is practically unused. Therefore, the issue of automating polygon control according to various constraints and conditions set by the client's body characteristics, as well as the appropriate clothing style, has not been addressed yet.

The quality of body measurements could be improved in various ways — by relying on anthropometric knowledge (statistical body proportions depending on gender, age, race, etc.), using 3D body scanning tools, visualizing clothing on digitally created avatars, and employing various statistical methods.

Each approach has its own advantages and drawbacks. For example, the use of 3D body scanning tools is limited by the relatively low availability of such equipment, specific

requirements for the clothing worn by the measured person during scanning, and dependence on built-in (and unchangeable from outside) calculation algorithms.

In this study, statistical methods will be used, focusing on identifying potentially erroneous measurements before pattern making. The set of measurements for one person will be subjected to statistical comparison with historically accumulated measurements of other individuals (calculating mutual correlations). If the correlation coefficient is below a critical threshold, additional measurement verification is required. This approach can be applied in various ways—by controlling the historical data development of one individual, by comparing the mutual relationships of individual measurements, or by introducing additional indicators (such as lengths and areas of pattern lines, etc.).

Of course, correct body measurements still do not guarantee that the clothing made with the respective patterns will fit the specific individual. There may be an error in the pattern generation script, or perhaps the particular style is not suitable for the person's body type. Similarly, an incorrect result may arise due to issues in garment modeling or changes in the technical infrastructure.

In the case of individually designed patterns, the test objects are numerous and diverse — software used to construct individual clothing patterns (scripts), software in which scripts are executed (construction platform), interfaces with third-party solutions performing specific actions such as visualization, comparison, object placement, etc. Each aspect potentially requires different testing approaches. This publication describes an approach for assessing the quality of body measurements.

II. QUALITY OF BODY MEASUREMENTS

The essence of automated quality control approaches for individually tailored clothing models is as follows: Designers create clothing models, program the garment generation algorithm (create scripts), and present them to clients. From various models, the client chooses the most suitable one and submits their body measurements, which can be numerous — sometimes even more than 50. The software should generate high-quality garment patterns for the client's specific measurements.

To verify the quality of generated patterns (both in terms of measurement compliance and software functionality), in real life, the selected clothing pieces (models) should be sewn from the patterns and fitted to the body, including comparing them with mass-produced analog model garments. Unfortunately, this approach requires immense time and material resources, especially if the measurement sets of different individuals are not analyzed for their similarity, which would allow for the creation of equivalence classes of measurement sets and the possibility of sewing one sample from each equivalence class. However, even using equivalence classes does not solve the problem because the required number of samples needs to be multiplied by the

number of garment models, leading to the realization that a full quality check with conventional means is impossible.

In this study, it is proposed to perform an automated model verification by analyzing the mutual correlation of body measurements and additionally calculating integral measurements - the area and perimeter of the garment's element, which are closely related (high correlation coefficient). If these derived values differ from the values predicted by the statistical forecast, it indicates that the corresponding pattern algorithms need to be checked. This method can also reveal inaccuracies in the input of client measurements.

In this section, we will examine methods for obtaining body measurements and analyze the reliability of the obtained values. Two methods will be analyzed:

(1) automatic measurement determination from photographs and images obtained with the help of a 3D scanner;

(2) measurements taken by the clients themselves, followed by statistical analysis of the obtained values, which may show significantly different values from other clients, if such differences exist.

In the digital environment, measurements either need to be determined automatically (for example, from photographs or using 3D scanners) or rely on the client's (typically non-professional) self-measurements.

A. Anthropometry and 3D Scanning

Automated body measurement determination from 2D and 3D images has been widely researched, but unfortunately, the results are not widely used in practice. This is determined by various factors, such as the quality of the photographs, inaccurate posing, clothing worn during photography, limitations of image recognition algorithms, restrictions on the transmission of personal data, and others. Accurate and reliable recognition of body measurements is extremely important in various fields, such as fashion, healthcare, ergonomics, and virtual reality.

The ability to accurately perceive and analyze body measurements is crucial in personalized product design, optimizing suitability, and enhancing user experience. Traditional methods of obtaining body measurements often rely on manual measurement techniques performed by trained professionals. However, the introduction of digital technologies has paved the way for alternative approaches that can improve the accuracy, efficiency, and accessibility of measurements. Two notable technological areas that have significantly influenced body measurement recognition are 3D scanning and image recognition.

There are several works that provide insights into the latest techniques, algorithms, and challenges in image-based body measurement recognition [2], [3]:

1. an overview of various methods and approaches used for obtaining human body measurements from images [4],

2. the use of image processing methods for human body measurement and virtual clothing fitting including algorithms

and methods for obtaining body measurements from images. [5],

3. a review of image processing methods used for automatic human body measurement, including various image analysis techniques, extraction algorithms, and measurement evaluation methods [3],

4. the use of deep learning methods to estimate human body measurements from images, including the use of Convolutional Neural Networks (CNN) and other deep learning architectures, to achieve accurate and stable measurement estimation [6].

Anthropometry, the measurement of human body dimensions, plays a crucial role in various fields such as ergonomics, clothing design, healthcare, and biometrics. Thanks to technological advancements, 3D body scanning has emerged as a powerful tool for capturing precise body measurements, offering a comprehensive and accurate alternative to traditional measurement method.

The dataset "IEEE IC 3DBP" [7] provides researchers with a valuable resource for comparative analysis and evaluation of various anthropometric methods in 3D body scanning. The study [3] compares various anthropometric measurement methods based on 3D body scanning. [8], [4] focuses on developing a population-specific anthropometric model based on 3D body scanning data. The findings emphasize the importance of considering population-specific anthropometric analysis variations for applications such as clothing design, ergonomics, and product development.

[5] investigates the use of machine learning algorithms in anthropometric analysis using 3D body scanning [9]. The study examines the application of machine learning models for automated measurement extraction, body segmentation, and anthropometric measurement prediction. The obtained data indicate the potential of machine learning methods to improve the efficiency and accuracy of anthropometric analysis in image datasets.

[3], [10], [11] proposes a comparative analysis of various 3D body scanning technologies and sensor technologies. It evaluates the performance, resolution, and accuracy of different scanning techniques such as structured light scanning, laser scanning, and depth sensing. The aim of [12], [13] is to confirm the accuracy and reliability of 3D body scanning measurements by comparing them with traditional anthropometric methods. The study provides a comparison between measurements obtained from 3D body scanning and manual measurements taken using calipers and measuring tapes. The collected data are used to validate the reliability and practicality of 3D body scanning as a reliable measurement method.

Unfortunately, practical pattern construction systems utilize these technologies to a very limited extent due to several unresolved issues. These include deficiencies in technical infrastructure and professional specialists, ethical considerations regarding client scanning, and other related factors.

B. Data Quality Assessment Using Correlation Methods

The focus of this study is on the quality analysis of measurements taken by the clients themselves or by their trusted persons. Several issues are observed when conducting these measurements:

1. measurements and techniques of taking them may vary depending on the pattern-making method. To standardize the process, visual materials and instructional videos can be used to guide clients on how to take accurate measurements for the specific method used in each case. However, practical experience shows that many clients tend to ignore instructions, either due to impatience or overconfidence in their skills, resulting in incorrect measurements,
2. individuals without sewing experience are unable to accurately measure the body, resulting in measurements that are too tight or too loose in the wrong places,
3. individuals cannot measure certain body dimensions themselves, thus they are forced to rely on assistance from others, which causes stress and additional errors.

If the obtained measurements are incorrect (they do not correspond to the specific body), the individual will end up with ill-fitting clothing pieces, even if the pattern construction algorithm is flawless.

If suspicious sets of measurements, which could arise due to erroneous actions, could be identified automatically, it would be possible to reduce the risk of creating inappropriate patterns. One method for identifying problematic sets of measurements could be to establish a mandatory relationship between measurement definitions and activating control mechanisms at the time of measurement registration. By using these relationships, clear errors could be filtered out, such as an impossible scenario where a woman's bust circumference is smaller than the underbust circumference. However, this approach does not help statistically identify combinations of measurements that are unlikely, as human bodies vary significantly. It is almost impossible to find universal measurement relationships based solely on experience. In this study, correlations of measurements could be analyzed using regression analysis and the capabilities of artificial intelligence on historically accumulated sets of body measurements.

In the human body, a series of measurements exhibit high correlation. For example, both arms or both legs are usually of practically equal length. The available body measurements can be used to calculate the mutual correlation for all pairs of measurements, and it can be observed that for some pairs of measurements, the correlation falls within the range of 0.8-0.9, while for others, it exceeds 0.9, sometimes even surpassing 0.97. For measurement pairs with a correlation above 0.9, it is advisable to create scatterplots of measurements for many clients and identify significant differences or "outliers". If outliers are detected, it indicates that the specific body data may be erroneous, but it does not

justify concluding that the data is incorrect; in such cases, the user should be informed that there may be issues with the data. Additionally, more complex regression analysis of measurements is necessary, considering multiple measurements as independent variables and determining the value of the dependent measurement using regression methods (such as linear regression, etc.).

C. Data Quality Assessment using Total Data Quality Methods

In this section, we will examine the extent to which the automated pattern-generating system can leverage the insights of the Total Data Quality Theory (TDQM).

In 2001, Redman proposed the following definition of data quality: "Data quality is the degree to which data satisfies the specific needs of a given customer." [14]. ISO 9000:2015 [15] provides the following definition of the concept of quality: "Quality is the degree to which consumer needs are satisfied; it represents all the features and characteristics of a product or service that meet customer demand." Quality is therefore a multidimensional concept and encompasses all aspects of how well data align with their purpose.

The notion from quality theory that data quality is relative and dependent on data usage is one of the central principles. When formalizing data quality requirements, it must be considered that they vary from one use case to another. For the same data, different requirements must be defined depending on its usage - distinct data quality specifications must be formulated. Data quality that satisfies all data uses, or in other words, "absolute" quality, is a goal to strive for but is rarely achieved.

The TDQM studies [16] explore a wide spectrum of data and information quality dimensions. As early as 1996, Wang and Strong [17] proposed 15 data quality dimensions, dividing them into 4 groups. In 2001, Redman [14] introduced 51 data quality dimensions, categorizing them into 9 quality groups. In 2013, the Data Management Association International UK Working Group [18] reduced the number of dimensions to six, thereby avoiding an overabundance of dimensions. This approach to data quality monitoring is currently used by the European Statistical Office - Eurostat. Six data quality dimensions are proposed, not tied to specific applications.: Completeness, Uniqueness, Timeliness, Validity, Accuracy, and Consistency.

Data quality is often associated with how accurate the data is. However, data quality is more than just accurate data. The dimensions are not tied to specific data applications, so they should be considered as universal requirements applicable to all applications where the substantive meaning of the dimensions is specified.

Previous research in the field of data quality has not resulted in a universally accepted common theory among researchers. The imprecision of the concept of dimensions serves as an obstacle to this. Most theoretical research is characterized by a wide range of data quality dimensions. Sometimes the number of data quality dimensions is not only too large (ranging from several dimensions to several tens

with a list of additional criteria specified for each dimension [19], but also the difference between some of them is almost imperceptible.

In various proposals, the same attribute is often used to indicate semantically different dimensions and vice versa [20]. The main issue is that the precise meaning of each dimension is still being discussed, and there is no consensus on its meaning and how it should be evaluated. According to Batini and Scannapieco [20], theoretical studies on data quality have not yet provided a unified system of data quality concepts. Some authors [21] propose another solution for determining and evaluating data quality. The proposed data object-oriented quality model consists of three main components:

1. data object defines the data whose quality needs to be analyzed,
2. the data quality specification defines the conditions that must be met for the data to be considered qualitative,
3. the quality assessment process.

If data quality is associated with data objects without linking it to the concept of dimensions, it is possible to define an unlimited number of data objects with various structures. Depending on the use case, different data quality requirements can be formulated for the same data object. Therefore, the proposed solution corresponds to the relative nature of data quality and can be applied to quality management in pattern generation.

III. DATA QUALITY EVALUATION AND IMPROVEMENT

To demonstrate how linear regression methods can be applied to improve the quality of body measurement data, this section will describe step by step the process that was conducted in the study using a real dataset of body measurements.

A. Data Selection

Research question (III.A): is there available data to execute a data quality assessment using correlation methods (refer II.B)?

The authors utilized a dataset of body measurements collected as part of a clothing modeling project. The data was accumulated over an extended period from various sources, including self-measurement and professional tailors. From this comprehensive dataset, 2,053 unique measurement profiles were selected. These profiles were collected between 2019 and 2022 as part of a project focused on developing algorithms for customized pattern making based on individual body measurements (M2M). To ensure consistency, all measurements adhered to specific techniques customized for the pattern-making method employed in the project. Tailors involved in the project were already familiar with these techniques, but clients received additional resources, including instructions and video demonstrations, to ensure proper self-measurement if applicable.

Three pattern constructions — bodice base, sleeve, and trousers – were generated for the unique measurement profiles. These constructions were chosen because they utilize the widest and most frequently used body measurements. Only the profiles with successfully generated (without errors) all three patterns in ,svg format were included in the next selection round. Errors typically occurred due to missing crucial measurements, but in rare cases, geometrical errors indicated illogical relationships within the measurements in a particular set.

Answer on research question (III.A): the authors of the study had access to a set of data collected during a M2M project with 2053 unique profiles. Additionally, they used (refer III.F) a dataset derived from Ansur II [24].

B. Data Classification

Research question (III.B): how to identify measurement profiles that are qualified for data quality assessment using correlation methods (refer II.B)?

The generated .svg files were manually reviewed, and the results were recorded in a table (Table I), noting the profile identifier, credibility (classification of measurements as determined by specialists), and measurements’ source. Only profiles representing women’s measurements were retained.

TABLE I.
CLASSIFICATION OF PROFILES BASED ON THEIR CREDIBILITY

1046	suspicious	measured by customer
1048	good	measured in house
1050	good	measured by customer
1051	suspicious	measured by customer
1053	good	measured by customer
1057	suspicious	measured by customer
1061	suspicious	measured by customer
1073	bad	measured by customer
1075	suspicious	measured by customer
1076	suspicious	measured by customer
1077	suspicious	measured by customer
1078	bad	measured by customer
1082	bad	measured by customer

The credibility of measurements was determined based on the following criteria:

Good - The generated pattern base constructions visually appear good and reliable; considering the measurements, they are proportionate and appropriate for the represented body size.

Suspicious - The generated pattern base constructions visually appear suspicious, which may mean that some measurements are disproportionately large or small compared to others in the set. For example, a disproportionately long shoulder slope may indicate an incorrectly entered shoulder length (Plg). Suspicious measurements were recorded in the comments.

Bad - The generated pattern base constructions visually appear bad and unreliable (such bases have not been sewn or tested on real people).

The source of the measurements was determined primarily based on the profile name and existing experience:

Measured in-house - Measurements for specific individuals were taken by a tailoring professional.

Measured by customer - Measurements for specific individuals were taken by the client or a tester following the instructions and videos provided.

Not sure - It is not possible to determine the source of the measurements, or there is uncertainty about which group it belongs to.

To create a set of measurement profiles that adequately represent real product users, the customer-measured profiles were divided into groups named “Plusminus” representing main chest circumference (Gkra) intervals (Fig 1), which historically correspond to specific construction methods upon which the programmed pattern constructions are built:

- Plusminus1 (over Gkra = 63 cm),
- Plusminus1.5 (over Gkra = 79 cm),
- Plusminus2 (over Gkra = 95 cm),
- Plusminus2.5 (over Gkra = 111 cm),
- Plusminus3 (over Gkra = 127 cm),
- Plusminus3.5 (over Gkra = 143 cm).




Fig 1. Division of customer profiles by Gkra.

Answer on research question (III.B): profiles where qualified by their origin and by evaluating the generated pattern base constructions for those datasets.

C. Acquiring the Minimum Data Set

Research question (III.C): how to evaluate the identified dataset selected in III.B regarding its completeness for data quality assessment using correlation methods (refer II.B)?

Based on the conclusions of the study [22], to accurately predict human measurements, 500 sets of individual measurements are sufficient. Taking this information into account, the number of profiles to be represented in each plusminus group was calculated (Table II), where first column represents the “Plusminus” group, the second column shows the percentage of this group within the sample dataset, and the third column indicates the number of individuals representing each group.

TABLE II. ESTIMATION OF THE REQUIRED NUMBER OF MEASUREMENTS

Row Labels	Sum of value_metric	Desired profile count in each group
Plusminus1	7.92%	40
Plusminus1_5	27.00%	135
Plusminus2	32.57%	163
Plusminus2_5	19.04%	95
Plusminus3	9.08%	45
Plusminus3_5	4.40%	22
Grand Total	100.00%	500

A new table (Table III) was created with 469 profiles, which includes profile ID, name generated based on the profile’s main bust circumference (Gkra) value, “Plusminus” group, credibility.

Credibility was assessed on a scale from 1 to 3 according to the following criteria:

1 - most reliable input data, as they were measured by specialists who have mastered the specific clothing construction method, and the generated clothing base set is recognized as good.

2 - less reliable input data, as although the generated clothing base set is recognized as good, it is not possible to determine whether the measurements were taken correctly, as they were taken by clients or people from test groups.

3 - least reliable input data. The generated clothing base set is considered suspicious because some measurements seemed disproportionate or the base itself visually unusual but not obviously poor.

TABLE III. UNIFIED DESCRIPTION OF PROFILE DATA

p_profiles_id	p_name	plusminus	credibility
190	66a	pm1	3
193	77a	pm1	3
261	86a	pm1_5	1
286	91a	pm1_5	1
288	84a	pm1_5	1
289	86b	pm1_5	1
291	101a	pm2	1
292	93a	pm1_5	1

The primary table was populated with all available data from credibility category 1. Unfortunately, only 230 profiles met these criteria. Therefore, the missing number of profiles in each “Plusminus” group of this category was identified (Table IV) to create an improved dataset in the future, targeting specialists who can measure individuals representing the missing size ranges. The last two columns

provide information of how many credible measurement profiles are needed in each group with or without including persons that are measured several times.

TABLE IV.
ESTIMATION OF AVAILABLE/MISSING MEASUREMENTS

Row Labels	Sum of value_metric	Desired profile count in each group	Profiles to be measured (w. duplicates)	Profiles to be measured (wo. duplicates)
Plusminus1	7.92%		40	23
Plusminus1_5	27.00%		135	8
Plusminus2	32.57%		163	92
Plusminus2_5	19.04%		95	61
Plusminus3	9.08%		45	39
Plusminus3_5	4.40%		22	7
Grand Total	100.00%		500	270

The missing table was filled with data from categories 2 and 3. Even including these less reliable data, the target of 500 profiles in the dataset has not been achieved (Table V).

TABLE V.
MISSING MEASUREMENTS PER PLUS-MINUS GROUP

	Plusminus1 (Gkra >= 63 cm)	Plusminus1.5 (Gkra >= 79 cm)	Plusminus2 (Gkra >= 95 cm)	Plusminus2.5 (Gkra >= 111 cm)	Plusminus3 (Gkra >= 127 cm)	Plusminus3.5 (Gkra >= 143 cm)
count right now	24	135	163	95	38	14
still needed	40	135	163	95	45	22
difference	-16	0	0	0	-7	-8

Interpreting the upcoming results, it's important to note that the missing 31 sets, needed to reach the minimum dataset size of 500 profiles, are spread across both very small and very large body size categories. In contrast, there were sufficient measurements in the average size groups.

To leverage existing resources, the authors opted to analyze readily available data for this initial phase. This would result in the development of a validation method for new measurements using an incomplete dataset. However, the method's design allows for seamless integration of a more robust dataset in the future, facilitating its adaptation to improved information.

Answer on research question (III.C): the amount of selected profiles is sufficient for mid-sized groups but insufficient for small and larger sized groups. The sample set is large enough to develop a system for validating measurement accuracy, but it is insufficient for making assumptions about underrepresented groups.

D. Determining Correlation between Measurements

Research question (III.D): which measurements in the selected dataset have high enough correlations with others to be forecasted and to help detect significant outliers, indicating poor data quality?

Using the prepared dataset of 469 profiles, correlation coefficients between measurements were determined (Fig 2).




Fig 2. Correlation coefficients between measurements

37 of the most used measurements in the chosen pattern making method were selected. For each pair of measurements, an interactive scatterplot was created in RStudio [26], marking potential outliers located outside two standard deviations from the specific dataset's mean. Examples are in Fig 3 and Fig 4.




Fig 3. Correlation of Gkra and Va (0.951254392293655)




Fig 4. Correlation of Ag and Bg (0.8286307496543)

27 measurements show a high enough correlation (coefficient over 0.7) with others. Additional 4 measurements have moderate correlation (coefficient of $0.5 < |r| \leq 0.7$), the rest have low or negligible correlation.

For pairs of measurements exhibiting a correlation coefficient of at least 0.7, profile IDs and reliability were recorded. Since one of the authors is a tailor and was involved in measuring process of bodies from data set, it was easier to determine nature of the outliers in credibility group 1 and it was assumed that they most likely represent natural variation in body measurements, while outliers with profiles representing categories 2 and 3 could also be due to incorrectly taken measurements or input errors.

Results were documented (Table VI). The column "outlier_percent" indicates the percentage of outliers marked with credibility categories 2 and 3. The column "could be true outlier" records outliers that are distinctly far from the rest of the points in the scatter plot. Outliers with credibility category 3 predominated in this category.

TABLE VI.
ANALYSIS OF OUTLIERS

Measurement_1	Measurement_2	Correlation_Coefficient	outlier_percent	2 un 3 credibility	1 credibility	could be true outlier
Gkra	Zkra	0.97	91	1083, 2545, 2613, 2956, 3389, 3531, 190, 193, 2709, 2781	368	307, 347, 381, 381, 2236, 2238, 2240, 2241, 2246, 2316, 2350, 2382, 2974, 3346, 3409
Ga1	Ga2	0.95	66	3301, 919, 738, 3505		2982 extreme, 3145 and 2709 close to extreme, 3055, 507, 1076 on the fence
Va	Zkra	0.95	100	998, 3281, 190, 507, 930, 2600	2250, 2252, 2315, 2895	
Krau	Zkri	0.95	60	930, 2600	2895	998
Gkra	Va	0.95	100	193		
Bg	PrB	0.94	97	a lot	3415	
Ciga	Pca	0.93	43	3475, 193, 475, 3282, 3367, 2600, 885,	324, 335, 2208,	324
Ikra	Pca	0.92	100	3040		3282 very prominent outlier

Summarizing the results of the outliers' analysis (Table VII with pie chart within), it is evident that the majority of outliers occurred in profiles with low reliability. Although seemingly trivial, this aspect is significant for further data utilization.

TABLE VII.

DISTRIBUTION OF OUTLIERS ACCORDING TO THEIR CREDIBILITY

unique id's	frequency	credibility
3341	26	3
2982	16	3
2823	13	3
3145	12	2
2709	12	3
324	9	1
2556	9	3
3282	8	2
1046	8	3
3367	7	3
412	6	3
922	6	3
3092	6	3
2600	5	3
2327	4	1
3409	4	1
2921	4	2
2934	4	2
834	4	3
193	3	3
507	3	3

Credibility 1:	27
Credibility 2:	55
Credibility 3:	160

Answer on research question (III.D): using the given set of data, 27 out of 37 measurements show a high correlation (coefficient over 0.7) to be used in the initial outlier analysis.

E. Predicting Measurements from Existing Data

Research question (III.E): with what accuracy can each of the 37 measurements be predicted?

Profiles with significant outliers were excluded from the data set which was split into training and testing sets.

The training set contained 80% of the data, the remaining 20% of data was reserved for testing. For each measurement a predictive model was trained. Later the model was used on testing data set, and the results were compared.

In the initial study, a linear regression model was used due to its simplicity, interpretability, and the natural linear relationships among body measurements. A previous study [22] and [23] also recommended using linear regression for estimating body dimensions, noting that it tends to be more stable compared to tree-based models. The 'train()' function in R was utilized because it offers a variety of parameter options. For instance, 'preProcess = c("scale", "center")' standardizes the data by giving it zero mean and unit variance, and 'trControl = trainControl(method = "LOOCV")' specifies Leave-One-Out Cross-Validation (LOOCV) for model training. LOOCV is a robust cross-validation method where each observation is used once as a test set, with the rest serving as the training set, which helps in thoroughly evaluating the model's performance.

Results for the main chest circumference predictions compared to actual measurements are pictured in scatterplot below (Fig 5). There is a linear regression line fitted over all data points and ellipses with confidence level of 95% for each group of credibility. As expected, the data points that are the most credible (credibility 1, red) are closer to regression line and therefore predicted better. Data obtained from clients (credibility 2, green) tend to be more scattered away from the regression line, particularly for the least reliable entries (credibility 3, blue).




Fig 5. Gkra vs. predicted Gkra

Based on the predictive model's performance, body measurements can be predicted with an accuracy of +/- 1.7 cm for the middle 68% of the data within the testing set. Furthermore, within the middle 95% of the data, the accuracy extends to +/- 4 cm. To provide context, the average bust circumference measurement is 95 cm. Therefore, this level of accuracy ensures errors are limited to less than 3.8% of this particular measurement itself. This represents a significant improvement over the previous approach used in validating measurements before generating M2M patterns, which relied mainly on detecting unlikely scenarios (e.g., an ankle circumference exceeding a knee circumference), empirical formulas based on only a few other measurements and minimal and maximal values of each measurement across the population. Although the prediction of this trained model does not meet the goal of +/- 1 cm it still can be leveraged to validate user-provided measurements to avoid gross user data entry errors. If a measurement falls outside the predicted bounds, a warning notification would be triggered, prompting the user to re-check the entered value.

The performance of the models was further analyzed for each measurement individually using metrics such as R² (Fig 9), adjusted R² (Fig 10), coefficient of variation (Fig 6), F-statistic (Fig 7), and its p-value (Fig 8). The results may show that the rankings of the best-performing models differ based on these metrics. For instance, model Ag (body height) might have a lower coefficient of variation (CV) for its residuals (Fig 6), indicating more consistent prediction errors, while model Gkra (bust circumference) might have a statistically more significant fit based on its F-statistic and p-value (Fig 7). However, the worst-performing models consistently show poor results across all metrics, indicating a clear distinction in predictability between the best and worst-performing measurements.

The measurement Pln (forward projection of the shoulder) demonstrated the least predictability using linear regression methods, as evidenced by a CV of 34.7%, an R² value of 0.210, an adjusted R² of 0.110, an F-statistic of 2.094, and a p-value of 0.0004806754. This lack of predictability is likely due to the observational nature of the measurement, which makes it prone to misinterpretation.

This approach also struggles to accurately predict PlslL (shoulder slope right), PlslKr (shoulder slope left), and Trg (trapezoid muscle length), which is unsurprising given the potential imprecision of these measurements, even within the highest credibility data group. Alternative methods should be explored to achieve more accurate results for these measurements, particularly for PlslL and PlslKr, as shoulder slope significantly impacts garment fit.

In contrast, the remaining measurements appear to be fairly predictable, with an average CV of 4.3% and a median CV of 4%.




Fig 6. Coefficient of Variation for each measurement model




Fig 7. F statistic for each measurement model






Fig 8. P values for each measurement model

Fig 9. R^2 for each measurement modelFig 10. Adjusted R^2 for each measurement model

Answer on research question (III.E): the accuracy of predictions is within ± 4 cm for majority of measurements.

F. Predicting the Impact of Additional Data

Research question (III.F): what impact would the addition of new, high-quality data have on the model?

To ensure that increasing the amount of high-quality measurement data would improve prediction accuracy, a dataset derived from Ansur II[24] data was created. Ansur II[24] data set consists of measurements of 1986 females. For the initial attempt, all measurements matching ours were selected, and the rest were calculated. For example, in our dataset, the Clg measurement is from the waist level to the middle of the kneecap. In Ansur II[24], knee height is measured from the middle of the kneecap to the floor [25]. Therefore, to obtain a comparable measurement to ours, the formula "Clg = Ansur II Iliodristal height [25] (similar to our PrB) – Ansur II Knee Height, Midpatella [25]" was used.

In cases where deriving the necessary measurements was not possible, they were assumed from our dataset. For instance, the Papla measurement correlates most with Gkra and Ra1. Since the Gkra measurement technique matches between our dataset and Ansur II [23], the Papla measurement

for each person in the Ansur II [23] database was assigned based on the closest Gkra measurement in our dataset. This approach was used to introduce some natural variability as opposed to calculated values, which would be immediately captured by the regression model, leading to overfitting. These measurements were flagged to facilitate measuring the model's performance with and without them due to their questionable nature.

Linear regression models for each measurement were trained twice: once with all the calculated and assumed measurements, and once with only the original measurements. Models trained with original and added data had an average CV of 3.25% and a median CV of 3.2%. Models trained with fewer but only original data performed slightly better with an average CV of 3.08% and a median CV of 3.17%.

The results of the model performance using the dataset with more quality data showed an improvement in prediction precision, as evidenced by lower mean and median CVs compared to the results of the linear regression model of our dataset.

We must consider that the Ansur II dataset lacks information for the "Plusminus3" and "Plusminus3_5" groups. Consequently, profiles from these groups were temporarily excluded from our dataset, and the performances of our measurement models were reassessed. The CV decreased to a mean of 3.91% and a median of 3.79% which is still worse than the results of models trained with Ansur II[24] data set. In conclusion it's safe to say that if two datasets with similar variance is used to train the model the one with more quality data will provide better results.

Answer on research question (III.F): using only high-quality data to train the predictive models leads to more precise measurement predictions.

IV. PRACTICAL APPLICATIONS

The results presented here regarding the discussed body measurements and the calculated prediction accuracy are specific — they apply 1:1 only to the specifically used methodology for creating patterns and the specifically available sets of body measurements. For the specific M2M case, based on the conducted analyses, concrete tolerance intervals can be defined for 33 body measurements. When implemented in an application, these intervals generate warnings if an entered value for a body measurement is likely inconsistent with the other body measurements. This application thus signals to the user a likely, statistically justified data inconsistency, preventing avoidable errors in data collection. However, for 4 body measurements (Pln, PlsIKr, PlsIL, Pln) discussed in section III.E, no reasonably justifiable tolerance intervals can be calculated, and no corresponding warnings can be generated. As explained above, an improvement (i.e., a reduction of the corresponding tolerance intervals) would be achieved by expanding the available data set and repeating the calculations. It is credible — at least for the middle "Plusminus" groups — that the goal

of 1cm tolerance for most body measurements is achievable with a three-digit number of data sets.

The specifically calculated 33 tolerance intervals cannot, of course, be directly transferred to other application areas. However, if we abstract from the specifically calculated values and interpret the applied approach as the investigation result, further statements about practical applicability can be made. The approach can fundamentally be transferred to any set of body measurements (both similar but differently defined ones, as well as completely different anthropological dimensions, such as body weight, which was not available to us). Wherever a set of (correlated) body measurements is manually recorded, and the data quantities are at least in the three-digit range, the outlined approach can be applied. Due to the anthropologically given correlation of body measurement data, it is possible to determine statistically justifiable tolerance intervals for most body measurements, and thus a practical method for avoiding errors in data collection can be defined. The outlined method is therefore not only applicable to the field of M2M pattern making used here, but also to the clothing industry in general and possibly even in completely different contexts (e.g., in the medical field).

V. CONCLUSIONS

The obtained results can be summarized as follows:

1. Even with a relatively small dataset of measurements (in this case, 469 profiles), it is possible to effectively identify potential data quality issues as long as a clear data quality assessment methodology is applied (refer the answer on research question III.D).
2. Using linear regression methods, it is possible to develop and train models that can reliably predict certain body measurements with high accuracy (refer the answers on research questions III.D and III.E).
3. The accuracy of measurement predictions could have been improved if more data from small and very large body sizes had been available (refer the answers on research questions III.C, III.D, III.E and III.F).
4. To get closer to the desired accuracy of ± 1 cm, it is essential to obtain more reliable data measured by professionals with a credibility rating of 1. The current dataset is sufficient to significantly improve the existing solution for measurement validation, which relies on the minimal and maximal values of each measurement across the population and empirical formulas. However, a predictive model based on the existing data would primarily identify gross errors and mistypes. More reliable data, beyond the desired 500 profiles, should be obtained to test the model's accuracy and determine the percentage of accurate and inaccurate predictions available (refer the answers

on research questions III.B, III.C, III.D, III.E and III.F).

The research is planned to be continued in two directions. Regarding the validation of body measurements, we intend to employ additional statistical methods for validation and compare them in terms of achievable statistical quality and practical applicability. Similar approaches can be found, among others, in [22].

Furthermore, we aim to investigate whether a similar use of statistical methods can also be employed to test M2M implementations (custom-made patterns). Assuming there is a pattern program that has been validated for N individuals, the question arises whether potentially faulty patterns for additional individuals can be identified without actually sewing and trying them on — the currently only practical way to validate an M2M pattern. Given the potentially avoidable costs involved, a machine-based, statistically driven method for identifying potential errors would be of very high practical relevance.

REFERENCES

- [1] J. Bicevskis, E. Diebelis, Z. Bicevska, A. Neimanis, "Regression Testing: Test Cases for Graphical Images." In Joint Proceedings of Baltic DB&IS'2022 Doctoral Consortium and Forum. <https://ceur-ws.org/Vol-3158/paper6.pdf>
- [2] Ruiz N., Bueno M.B., Bolkart T., Arora, Lin M., Romero J., Bala R., "Human body measurement estimation with adversarial augmentation. International." Conference on 3D Vision, 2022 <https://arxiv.org/abs/2210.05667>
- [3] Bartol K., Bojanić D., Petković T., Pribanić T., A "Review of Body Measurement Using 3D Scanning." IEEE Access, DOI: 10.1109/ACCESS.2021.3076595, 2021.
- [4] Kuribayashi M., Nakai K., Funabiki N., "Image-Based Virtual Try-on System With Clothing-Size Adjustment." DOI: 10.48550/arXiv.2302.14197, 2023.
- [5] Pleuss J.D., Talty K., Morse S., Kuiper P., Scioletti M., Heymsfield S.B., Thomas D.M., "A machine learning approach relating 3D body scans to body composition in humans." Eur J Clin Nutr. 2019 Feb; 73(2): pp. 200–208, published online 2018 Oct 12. doi: 10.1038/s41430-018-0337-1
- [6] Ashmawi S., Alharbi M., Almaghrabi A., Alhothali A., "Fitme: Body Measurement Estimations using Machine Learning Method." Procedia Computer Science. Volume 163, pp. 209-217, 2019. <https://doi.org/10.1016/j.procs.2019.12.102>.
- [7] "IEEE IC 3DBP" <https://iee-dataport.org/open-access/dataset-ieee-ic-3dbp-comparative-analysis-anthropometric-methods>.
- [8] Lu J., Wang M.J., "Automated anthropometric data collection using 3D whole body scanners." DBPL, Expert Systems with Applications 35(1-2):407-414, July 2008. DOI: 10.1016/j.eswa.2007.07.008
- [9] Liu X., Wu Y., Wu H., "Machine Learning Enabled 3D Body Measurement Estimation Using Hybrid Feature Selection and Bayesian Search." Appl. Sci. 2022, 12(14), 7253; <https://doi.org/10.3390/app12147253>.
- [10] Kus A., Unver E., Taylor A., "A Comparative Study of 3D Scanning in Engineering, Product and Transport Design and Fashion Design Education." Computer Applications in Engineering Education 17(3), pp. 263 –271, September 2009 DOI: 10.1002/cae.20213
- [11] Seifert, E., Griffin, L., "Comparison and Validation of Traditional and 3D Scanning Anthropometric Methods to Measure the Hand. Paper presented at 11th Int. Conference and Exhibition on 3D Body Scanning and Processing Technologies. <https://doi.org/10.15221/20.41>, 2020.
- [12] Skorvankova, D., Riečický, A., Madaras, M., "Automatic Estimation of Anthropometric Human Body Measurements." 17th International Conference on Computer Vision Theory and Applications. (2021) DOI:

- 10.5220/0010878100003124,
<https://www.scitepress.org/PublishedPapers/2022/108781/108781.pdf>
- [13] Rumbo-Rodríguez, L., Sánchez-SanSegundo, M., Ferrer-Cascales, R., García-D'Urso, N., Hurtado-Sánchez, J.A., Zaragoza-Martí, A., "Comparison of Body Scanner and Manual Anthropometric Measurements of Body Shape: A Systematic Review." *Int J Environ Res Public Health*. 2021 Jun 8;18(12):6213. doi: 10.3390/ijerph18126213
- [14] T.C. Redman, "Data Quality. The Field Guide", Digital Press, 2001.
- [15] Quality management systems.
<https://www.iso.org/standard/62085.html>
- [16] European statistics Code of Practice — revised edition 2017.
<https://ec.europa.eu/eurostat/web/products-catalogues/-/KS-02-18-142>
- [17] R. Y., Wang, D. M., Strong, "Beyond Accuracy: What Data Quality Means to Data Consumers", *Journal of Management Information Systems*, Springer, Vol.12., No.4, pp. 5-34, 1996.
- [18] DAMA UK. <https://www.dama-uk.org/>
- [19] A. Caro, C. Calero, M. Piattini, "A Portal Data Quality Model for Users And Developers", In *ICIQ*, pp. 462-476, 2007
- [20] C. Batini, M. Scannapieco, "Methodologies for Information Quality Assessment and Improvement. Data and Information Quality Dimensions, Principles and Techniques", Springer International Publishing Switzerland, 2016.
- [21] J. Bicevskis, Z. Bicevska, A. Nikiforova, I. Oditis, "An Approach to Data Quality Evaluation". In *Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pp. 196-201, IEEE, 2018.
- [22] Meyer, P., Birregah, B., Beuseroy, P. et al., "Missing body measurements prediction in fashion industry: a comparative approach." *Fash Text* 10, 37 (2023). <https://doi.org/10.1186/s40691-023-00357-5>
- [23] Janis Bicevskis, Edgars Diebelis, Zane Bicevska, Ivo Oditis, Girts Karnitis, Oskars Ozols. "Assessing the Accuracy of Body Measurements through Regression Analysis." In *18th Conference on Computer Science and Intelligence Systems*, September 17–20, 2023. Warsaw, Poland : Position Papers Vol. 36, p.35-41. DOI: <http://dx.doi.org/10.15439/978-83-969601-1-5>.
- [24] The Anthropometric Survey of US Army Personnel (ANSUR 2) Female data files, 2012. <https://www.openlab.psu.edu/ansur2/>
- [25] Hotzman, J., Gordon, C.C., Bradtmiller, B., Corner, B.D., Mucher, M., Kristensen, S., Paquette, S., and Blackwell, C.L. 2011, "Measurer's Handbook: US Army and Marine Corps Anthropometric Surveys, 2010-2011," Report No. NATICK/TR-11/017.
- [26] RDocumentation. <https://www.rdocumentation.org/packages/caret/versions/4.47/topics/train>

Using a Textual DSL With Live Graphical Feedback to Improve the CPS' Design Workflow of Hardware Engineers

Twan Bolwerk

Philips in Best, The Netherlands &
Radboud University in Nijmegen,
The Netherlands

Marco Alonso

Philips in Best, The Netherlands

Mathijs Schuts

Philips in Best, The Netherlands &
Radboud University in Nijmegen,
The Netherlands

Abstract—Cyber-Physical Systems are designed and developed using multi-disciplinary teams. Handovers from one discipline to another often occur using text documents written in natural language, which can be imprecise, ambiguous, and lead to errors. To improve this situation, we created a textual Domain Specific Language with live graphical feedback to enhance the handover between mechanical and mechatronic engineers working on medical robots at Philips IGT. The Domain Specific Language formalizes the system description and provides immediate live graphical feedback to prevent mistakes from being made, such as editing the wrong physical parts and by visualizing the differences of two versions of a system. In addition, our approach leverages multiple industry standards and it enables bi-directional navigation between languages.

I. INTRODUCTION

A CYBER-PHYSICAL System (CPS) [4] is a complex system composed of both hardware and software components. These systems are designed and developed with multi-disciplinary teams. Often, a CPS consists of moving parts such as in the case of robots, cars, airplanes, etc. For the hardware component development, mechanical and mechatronics engineers are involved. The mechanical engineer creates 3D models of the physical components using a Computer Aided Design (CAD) tool [27] and performs measurements, i.e., on weight and tolerances. The mechatronics engineer makes these physical systems move by creating control solutions using tools such as Matlab and Simulink [21]. Both disciplines use their own specialized software tools. Currently, the workflow involves manually written documents that are used to handover designs and measurement information from the mechanical engineer to the mechatronics engineer. Due to the informal nature of these documents, they can be imprecise, ambiguous, and prone to errors. Additionally, changes between document versions may go unnoticed.

At Philips IGT, we create interventional X-ray systems such as the Azurion system in Figure 1, which are used for minimally invasive procedures. These large medical robots feature motorized moving parts that can be operated using joysticks [24].

In this paper, we present an improved workflow for the development of these CPSs. After a mechanical component




Fig. 1. Interventional X-ray system

has been modelled using a CAD tool, the mechanical engineer can export the 3D model in the Unified Robot Description Format (URDF) [17], which is based on eXtensible Markup Language (XML) [7]. However, one downside of URDF is that weights and tolerances are not included, and cannot be added. Additionally, it lacks an import mechanism for reusing components across similar robots. These limitations can be addressed using the XML macro language (Xacro) [2], which requires manual editing to add weights and tolerances. Both formats are in XML, which is not a user-friendly way of editing. Furthermore, we place these files in a version controlled system, but merging XML-based files is challenging.

The first author of this paper created a textual Domain Specific Language (DSL) [11] called Geometry Specific Language (GSL) or GeometrySL. The language extends Xacro but is not based on XML. Instances of this language are placed in a version controlled system and handed over from mechanical engineers to mechatronics engineers. By using formal GSL files instead of informal documents, we reduce the likelihood of errors due to handovers. The GSL provides immediate live graphical feedback when editing the textual instance, showing which part is being edited within the robot's 3D model. It also has facilities for graphically comparing two versions of a robot, highlighting parts that are different in a 3D model. Additionally, it supports bi-directional navigation from a graphical part to the corresponding DSL fragments and vice versa. This allows seamless navigation between graphical

views, URDF instances, Xacro instances, GSL instances, and back again using shortcuts.

To the best of our knowledge, the novelty of this research lies in the creation of the GSL, a DSL that defines how differences between robot representations are visualized. By leveraging multiple languages, including industry standards like Xacro and URDF, this approach enables bi-directional navigation and offers a unique method for visualizing differences in robot descriptions.

The paper is organized as follows. In Section II, we provide an overview of related work. We describe the current and proposed workflows in more detail in Section III. The GSL, Xacro and URDF languages are presented in Section IV. Section V describes the design of the tool. The resulting tool is shown in Section VI. Discussion is in Section VII. In this section, we also discuss how our work is related to the work of others. And we conclude our paper in Section VIII.

II. RELATED WORK




Fig. 2. DSL categories from [26]

Shen et al. [26] categorized recent studies on DSL based on three concerns: concrete syntax, abstract syntax and semantics. They analyzed the parsing and mapping strategies of these studies to classify them into categories such as external/internal, textual/graphical, modeling/visualizing/embedding. This study aims to address research gaps in DSL categorization. The vertical axis of Figure 2 lists literature references while horizontal axis list the following categories:

- **External (Ext):** Standalone languages with their own syntax and grammar, distinct from any host language, providing specific solutions within a particular domain.
- **Internal (Int):** Embedded within an existing general-purpose programming language, leveraging the host language’s syntax and features to implement domain specific constructs.
- **Textual (Tex):** Use text-based syntax, similar to traditional programming languages, designed for domain experts familiar with coding or scripting.
- **Graphical (Gra):** Employ visual representations such as diagrams and flowcharts to define domain specific constructs, useful for users preferring visual over textual representation.
- **Domain Specific Visual Language (DSVL):** A subset of graphical DSLs using specialized visual notations to represent domain concepts, facilitating understanding and communication among stakeholders.

- **Domain Specific Modeling Language (DSML):** Focus on creating models specific to a domain using specialized syntax and semantics, providing tools for simulation, validation, and code generation.
- **Domain Specific Embedded Language (DSEL):** A type of internal DSL embedding domain specific constructs within a host language, integrating domain specific functionality directly into general-purpose language code.

The figure features a red box highlighting the scarcity of research focused on Gra and Tex DSL combinations. While enabling live graphical feedback reduces cognitive load, improves collaboration and communication in engineering projects. This study bridges computing science and behavioral research domains to improve DSL design with live graphical feedback for better communication and collaboration in complex engineering projects.

To visualize property changes in language instances, we need to consider their visibility. Munzner’s book [19] recommends automatic highlighting with varied colors, shapes, or positions can emphasize distinctions between properties. The “pop-out” effect in Munzner’s book helps users spot differences quickly without focused attention.

In [14], Joshua Horowitz et al. define programming qualities. The focus is on immediate feedback (liveness), domain specific editing (richness) and composability. Composability enables the inclusion of external libraries or components, separating responsibilities over multiple sources. Programming tasks often require using multiple composed tools, so effects should be visible with minimal distraction and effort. Liveness and richness often fail to retain their composability according to Horowitz et al. conclusion. They identified a trend where interfaces lacking composability are standalone applications that offer limited utility in practice. This research explores the intersection of liveness, richness and composability by adhering to these qualities using familiarity with 3D graphical tools for hardware-engineers’ workflow improvement and tool intuitiveness enhancement.

Van Rozen et al. [34] recognized the need for program execution observation in traditional programming, which requires re-execution of updated source code from the beginning. This process is time-consuming and distracting when valuable states are lost or difficult to reproduce. To address these challenges, they propose a more fluid and live experience in programming using Textual Model Diff (TMDiff) [31]. Tmdiff uses two key techniques: origin tracking (tracing semantic model elements back to their defining source code) and text differencing (identifying corresponding model elements when aligned names have the same origin). The deltas found by TMDiff are converted into run-time edit operations, which can be applied atomically using rmpatch. Custom state migrations extend rmpatch to avoid information loss or invalid run-time states. Events like user interactions and changes in source code are recorded for undo functionality, persistent application state and back-in-time debugging. They evaluated existing methods (Xtext and EMFCompare) using Eclipse Modeling Framework (EMF) and found TMDiff’s scope-handling ability

more flexible. Their goal is to minimize distractions and preserve intermediate visual state for a smoother programming experience.

In [10], Cooper et al. present requirements and challenges to integrate graphical editors using Sirius within EMF. They note that while Sirius allows creation of custom modeling editors, it has a steep learning curve. To address this limitation, they propose five requirements for a hybrid textual-graphical workbench: syntax-aware editing, scoping and referencing, rename refactoring, error/warning marker display and accessibility to the textual model. These requirements aim to improve productivity, reduce errors and facilitate collaboration by providing seamless integration between graphical and textual models. Their proposed solution aligns with a case study on hybrid modelling workbenches.

A DSML for UML profiles was created by combining Papyrus and Xtext using EMF. One unique feature is shared storage base for both textual and graphical views, reducing synchronization efforts. The tool is tested by four scenarios (Create1, Modify1, Create2, Modify2) with experienced developers in the UML language. Results showed that creating elements and setting properties were faster in textual notation while constructing state machines was quicker in the graphical view. Renaming operations were faster in textual mode due to regex search and replace efficiency. The hybrid solution was superior in efficiency, doubling the speed in mentioned scenarios. This demonstrates potential of combined graphical-textual approach for DSMLs [1].

In game development, rapid adjustments of rules are made using tools like Machinations¹. Van Rozen et al. created a DSL called Micro-Machinations (MM) using the tool Rascal [30] to balance games. MM allows for direct visualization model editing, shortening feedback loops and reducing design iteration times by improving flexibility and adaptability. The Rascal Language WorkBench (LWB) with SPIN model checker [5] is used for analyzing MM, providing an IDE that reads textual MM and displays a visual model interactively. The MM Lib is embedded in the game itself to tackle interoperability, traceability and debugging challenges [29]. An immediate feedback loop greatly improves multi-disciplinary team collaboration in gaming domain similar to engineering domains.

Perez et al. [22] created DSVL using both textual and graphical views with ATOM tool². They followed meta-model centric approach where EBNF grammar was generated based on the meta-model, allowing decision to be made later whether to use graphical or textual syntax. Another issue is that produced Abstract Syntax Tree (AST) from parsing is not formally defined causing problems in integration with multi-view DSL proposed by them. They noticed that it is more natural to describe equations in a textual notation.

A more extensive version of this work can be found in [6].

¹<https://machinations.io/>
²<http://atom3.cs.mcgill.ca/>

III. WORKFLOW

In this section, we describe the current workflow of hardware engineers at Philips IGT and we propose a new improved workflow. The workflow involves two actors: mechanical engineers and mechatronics engineers.

A. Current workflow

The workflow process follows a waterfall approach, where mechanical engineers measure the system in the factory. These measurements are then communicated via various Office tools to mechatronics engineers. The tools used by these actors are depicted in Figure 3 and illustrate their interactions. The interactions between the tools can occur either automatically, with data being stored or transferred automatically between tools, or manually inputted by the user.

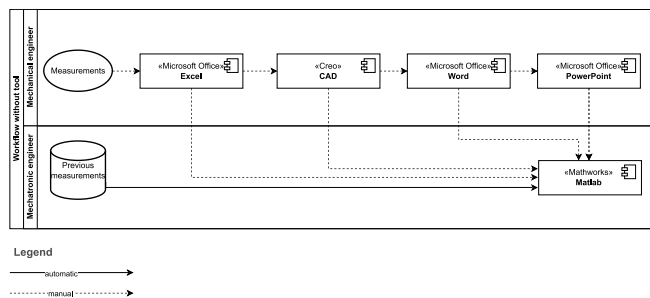


Fig. 3. Tools and actors.

We describe the actions per actor:

- Mechanical engineer. The mechanical engineer creates the hardware design of the system using the CREO³ tool for opening and editing 3D CAD files. They also measure the real-world properties of the system, perform calculations in Excel and document any changes compared to the original CAD model using Word. These updates are presented via PowerPoint [25].
- Mechatronic engineer. The mechatronic engineer manually compares previous measurements saved as Matlab instances with the changed properties recorded in CAD and Excel to determine if the system still meets the specifications. For example, they ensure that the system adheres to the predefined tolerances for each link⁴. These assessments are crucial for maintaining the system's performance and reliability, and are calculated using complex calculations and simulations in Matlab and Simulink [36].

In sum, the handover from the mechanical engineer to the mechatronics engineer currently relies on informal document-based communication. This can result in misunderstandings and potential errors due to missed changes or ambiguities.

³<https://www.ptc.com/en/products/creo/>

⁴A link in robotics refers to a rigid component that forms part of a robot's structure, connecting to other links through joints.

B. Proposed workflow

In the proposed workflow, all Microsoft Office tools are replaced by Domain Specific Modelling (DSM) [12] using a single Geometry Specification Language (GSL) with live graphical feedback. It replaces the Microsoft Office tools by a unified language capable of presenting changes, serving as documentation and evaluating mathematical expressions that can be used to describe and calculate properties.

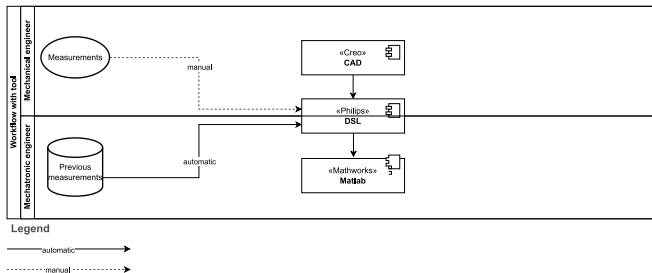


Fig. 4. Actors and tools workflow with DSL.

By eliminating the reliance on Microsoft Office tools and therefore multiple sources of truth, the complexity associated with using these tools is reduced. This can be observed in Figure 4 compared to Figure 3.

A CAD file can be converted into URDF which we are able to translate to our own DSL. The mechanical engineer can input measurements manually in an expressive manner similar to an Excel spreadsheet but now in the GSL. Meanwhile, the mechatronic engineer can compare previous measurements stored in the single source of truth DSL which is stored within Philips' version control system. Using a textual DSL makes, i.e., merging branches easy.

C. Use cases

In this section, we describe two main use cases. One involves presenting the textual DSL in a graphical manner and another focuses on visualizing differences between two versions of the DSL in a graphical representation.

1) *Present*: The first use case demonstrates effective communication of measurement changes between mechanical and mechatronic engineers. It serves as a visual representation tool, enabling the presentation of changes through visualization. By hovering over the textual representation using the cursor as a pointer, the view automatically focuses on the specific physical link of the robot being hovered over. This visualization feature enhances communication by clearly highlighting and displaying the changes made to each physical link of the robot. It provides a seamless and intuitive way for stakeholders to understand and interpret updated measurements for every physical link in the robot.

The mechanical engineer can efficiently navigate, inspect, and present the robot's hardware properties using the tool's textual and graphical representations. This use case showcases how engineers can leverage the tool's features to communicate and demonstrate the robot's hardware-properties.

2) *Highlight*: The highlight use case offers an efficient solution for mechatronic engineers to automatically identify and highlight changed physical links between robot or component versions. This internal scenario eliminates the need for manual input of values into Matlab, as it leverages automatic highlighting and live graphical feedback. Engineers can quickly and accurately pinpoint variations between link versions using this feature, reducing reliance on manual comparisons and potential errors.

IV. LANGUAGES

In this section, we introduce the Language WorkBench (LWB) used to implement our tool as needed in order to understand this paper. We describe the URDF, Xacro and GSL languages. Finally, we provide example instances of these three languages.

A. Language workbench

A LWB provide concepts and mechanisms to define language syntax, semantics, and code generation for a language. They facilitate model-driven engineering by allowing developers and domain experts to work with high-level abstractions that closely resemble the specific problem domain. This can lead to more efficient development processes, maintainable code, and closer collaboration between different stakeholders. LWBs bridge the gap between generalized programming languages and the unique requirements of specialized domains [8].

There are many language workbenches (LWBs) available. Both Eclipse Modeling Framework (EMF) and Rascal⁵ are used at Philips; however, Rascal provides support for Visual Studio Code (VS Code). Given that VS Code⁶ is the preferred tool at Philips and aligns with our previous experience using Rascal, we decided to use it for our tool.

Rascal [16] is a type-safe programming language featuring immutable data, built-in pattern matching, search, and relational calculus. It is a functional and procedural programming language with Java-like syntax. We introduce the language in this section as needed to understand the code fragments presented in this paper. The code snippets in Listing 1 are reused from [23].

```

1 loc l =
2   |file:///Users/kees/.bashrc|(100,20,<2,0>,<2,20>)
3
4 data Boolean = true() | false() | and(Boolean lhs,
5   Boolean rhs);
6 //extending Boolean with another constructor
7 data Boolean = or(Boolean lhs, Boolean rhs);
8 data Statement = \if(Expression c, Statement tt,
9   Statement ff);
10
11 for (int i <- [1 .. 5]) println(i);
12
13 str w = "world"
14 println("Hello, <w>!"); // prints "Hello, world!"

```

Listing 1. Rascal code fragments

⁵<https://www.rascal-mpl.org>

⁶<https://code.visualstudio.com>

Rascal has numerical types such as `int` and booleans, represented by `bool`. It supports polymorphic `lists`, `maps` for collections and `loc` for location constants.

Rascal also provides the following built-in functions for working with `maps`:

- 1) `Map` := A list that uses any kind of data as index (called keys), to store a value.
- 2) `rangeR` := Expects a map and returns a map of key, value pairs that match the values.
- 3) `range` := Returns a list of values.
- 4) `domain` := Returns a list of keys.

On Line 2 there is a constant `loc` that points to a file with the file scheme and selects the part on line 2 between the left margin and the 20th column. Locations are used to refer to files, store information extracted from files and help in referring back to source locations.

In Rascal, Algebraic Data Types (ADTs) can be user-defined with their constructor functions. The fragment on Lines 4-9 shows a declaration of an ADT for the representation of Boolean expressions using three constructors. The next line extends the same ADT by adding an alternative to the existing declaration. Reserved keywords are not permitted as names of algebraic constructors; hence, `if` is escaped with `\if` when used as the name of an algebraic constructor.

Control structures such as `for` can be used to iterate over a value, while `str` literals in Rascal are not delimited by line endings.

We utilized the following two Rascal libraries: `Salix` and `TypePal`. `Salix` is a library that facilitates the development of web-based GUI programs; it runs user code on the server side instead of client-side execution. The library employs the Model View Controller (MVC) pattern by sending HTML patches to the browser and interpreting messages from the browser on the server to update the view accordingly. `TypePal` is a typechecking and validator library for Rascal, designed to analyze and enforce type constraints, ensuring correct usage of types and detecting potential type-related errors in DSL instances. `TypePal` provides static type checking capabilities and can be used to improve the reliability and correctness of language instances [28].

B. URDF

Utilizing URDF offers numerous advantages, notably its capability to express a wide range of hardware properties, with the exception of tolerances. Moreover, URDF facilitates seamless conversion from formats like CAD files, streamlining the integration of engineering disciplines for our use case. Furthermore, the compatibility of URDF with visualization tools such as `RViz`⁷ underscores its ability in conveying essential information using a 3D graphical representation.

The downside of URDF is its inability to scale and its cumbersome XML format, which is difficult to maintain and causes issues in the archiving system when merging changes. As the model becomes more complex, the lack of reusable components results in larger file sizes.

⁷<http://wiki.ros.org/rviz>

C. Xacro

To address the issue of large URDF files, especially for complex 3D robot models like interventional X-ray systems, a solution based on modularization and composition using the Xacro language can be employed. Xacro provides a way to create modular and reusable components, making robot descriptions more manageable and organized.

Although Xacro simplifies URDF composition and introduces expression evaluation, it still relies on an XML format that is neither user-friendly nor intuitive, making it difficult to version control and prone to merging issues.

D. GeometrySL

To overcome the previously described limitations, we introduce GeometrySL (GSL), which extends XacroSL. GeometrySL encapsulates Xacro and converts it to a more user-friendly syntax, making it easier to archive and merge changes. It enabled the creation of a more intuitive syntax that supports custom property definitions (including tolerances) and provides enhanced visualization options. The aim is to offer a more intuitive, flexible language for 3D robot modeling, addressing the drawbacks associated with traditional XML-based URDF descriptions.

Creating GeometrySL adds flexibility in defining and designing custom syntax, making it more intuitive to use and easier to extend with new semantics and introduce custom properties, enhancing its expressiveness and adaptability for different use cases.

One of these new use cases is the integration of visual semantics. This allows mechanical engineers to describe desired appearances of views using a language that engineers can understand. This approach enables better collaboration by providing live graphical feedback.

E. Example instances

Next, we will describe three features using language instances—custom features, modular includes, and highlighting differences between two robot versions. Due to confidentiality concerns, we use Franka's Panda robot⁸ as an example instead of our interventional X-ray system.

```

1 robot {
2   link {
3     name="lbr_iiwa_link_0"
4     inertial {
5       origin := {
6         rpy="0 0 0"
7         xyz="-0.1 0 0.07"
8       }
9       mass := { value="0.2" }
10      tolerance := { value="200" }
11      inertia := {
12        ixx="0.05"
13        ixy="1"
14        ixz="0"
15        iyy="0.06"
16        iyz="0"
17        izz="0.03"
18      }

```

⁸https://support.franka.de/docs/franka_ros.html

```

19 }
20 visual {
21   origin := {
22     rpy="0 0 0"
23     xyz="0.2 0.1 0"
24   }
25   geometry {
26     mesh := { filename="meshes/link_0.stl" }
27   }
28 }
29 collision {
30   origin := {
31     rpy="0 0 0"
32     xyz="0 0 0"
33   }
34   geometry {
35     mesh := { filename="meshes/link_0.stl" }
36   }
37 }
38 }
39 }

```

Listing 2. GeometrySL instance example

Listing 2 shows an example of GeometrySL for the Panda robot, demonstrating how to define a link. A link has a name and various features such as inertial, visual, and collision properties. Inside the inertial feature, we added a custom property called “tolerance”. This property can be exported to Xacro but not to URDF. The instance also references STereo Lithography (STL) files, which is a format used to describe 3D objects using triangles.

```

1 robot {
2   name="lbr_iiwa"
3   xmlns:xacro="http://www.ros.org/wiki/xacro"
4   include "lbr_iiwa_link_0.gsl"
5   ...
6   include "lbr_iiwa_link_7.gsl"
7
8   include "lbr_iiwa_joint_1.gsl"
9   ...
10  include "lbr_iiwa_joint_7.gsl"
11 }

```

Listing 3. GeometrySL instance example for modular includes

In Listing 3, an instance of GeometrySL is shown that describes the Panda robot. It includes other instances of GeometrySL to describe links and joints of the Panda robot. A link is represented as a mesh, while a joint defines the relation between exactly two joints.

```

1 <?xml version="1.0"?>
2 <robot name="lbr_iiwa"
3   xmlns:xacro="http://www.ros.org/wiki/xacro" >
4 <xacro:property name="color" value="Green"/>
5 <xacro:property name="half" value="0.1"/>
6 <xacro:include filename
7   ="lbr_iiwa_link_0.gsl.xacro"/>
8 ...
9 <xacro:include filename
10  ="lbr_iiwa_link_7.gsl.xacro"/>
11
12 <xacro:include filename
13  ="lbr_iiwa_joint_1.gsl.xacro"/>
14 ...
15 <xacro:include filename
16  ="lbr_iiwa_joint_7.gsl.xacro"/>
17 </robot>

```

Listing 4. Xacro instance example for modular includes

Listing 4 shows an example of how to represent the same information from Listing 3 in the Xacro language. It uses XML syntax instead.

```

1 <?xml version="1.0" ?>
2 <robot name="lbr_iiwa">
3   <link name="lbr_iiwa_link_0">
4     <inertial>
5       <origin rpy="0 0 0" xyz="-0.1 0 0.07"/>
6       <mass value="0.2"/>
7       <inertia ixx="0.05" ixy="1" ixz="0"
8         iyy="0.06" iyz="0" izz="0.03"/>
9     </inertial>
10    <visual>
11      <origin rpy="0 0 0" xyz="0.2 0.1 0"/>
12      <geometry>
13        <mesh filename="meshes/link_0.stl"/>
14      </geometry>
15      <material name="Grey"/>
16    </visual>
17    <collision>
18      <origin rpy="0 0 0" xyz="0 0 0"/>
19      <geometry>
20        <mesh filename="meshes/link_0.stl"/>
21      </geometry>
22    </collision>
23  </link>
24  ...
25  <link name="lbr_iiwa_link_7">
26    ...
27  </link>
28  <joint name="lbr_iiwa_joint_1" type="revolute">
29    <parent link="lbr_iiwa_link_0"/>
30    <child link="lbr_iiwa_link_1"/>
31    <origin rpy="0 0 0" xyz="0 0 0.1575"/>
32    <axis xyz="1 0 1"/>
33    <limit effort="300" lower="-2.96705972839"
34      upper="2.96705972839" velocity="10"/>
35    <dynamics damping="0.5"/>
36  </joint>
37  ...
38  <joint name="lbr_iiwa_joint_7" type="revolute">
39    ...
40  </joint>
41 </robot>

```

Listing 5. URDF instance example for expanded Xacro instance

Listing 5 shows an example of how to represent the same information as Listing 4 using URDF syntax instead. It includes expanded information from Listing 2, but does not include the “tolerance” property that was present in GeometrySL and Xacro.

```

1 highlight
2 robot "robot_v1.gsl"
3 difference
4 robot "robot_v2.gsl"

```

Listing 6. GeometrySL instance example for highlighting differences for two versions of a robot

Listing 6 demonstrates a language instance of GeometrySL that visually shows the differences between two versions of the robot.

V. DESIGN

In this section, we describe how we realized the use cases presented in Section III.

A. Conversion tool

One of the benefits of using URDF, a widely adopted and standardized format is that other tools often have conversion features. For instance, the Blender tool⁹ can be utilized as an intermediary that enables conversion from CAD export files into URDF. Blender is open-source, has free licensing, and it has widespread community support. This not only makes Blender cost-effective but also ensures that the tool is consistently updated and improved by a global community of developers.

B. VS Code URDF viewer

The VS Code URDF viewer¹⁰ was utilized as the starting point, which incorporates BabylonJS¹¹, a JavaScript 3D graphics library. This library can load STL files and assemble them using URDF. However, it lacks certain features such as comparing changed properties, highlighting differences, bi-directional navigation and side-by-side views, and maintaining state after changes. Additionally, it requires the use of URDF, which has a non-user-friendly syntax and leads to large file sizes when designing systems like interventional X-ray systems. Nonetheless, despite these limitations, the VS Code URDF viewer serves as a beneficial starting point.

The visualization is extended with a context menu. The view's context menu enhances user experience by enabling the identification of links through right-click actions, revealing a menu that displays the link's name, depicted in Figure 6. This visual representation establishes a direct connection between the textual and visual physical links of the robot, improving cohesiveness and comprehension of both representations.

The sliders facilitate joint movement, enhancing the view by providing an interactive experience. While primarily focused on static properties, this feature can greatly improve the visualization of specific links. Furthermore, the “reload” and “auto” buttons play a vital role in updating the view with the latest changes made in the GSL textual editor, whether through manual input or automatic updates.

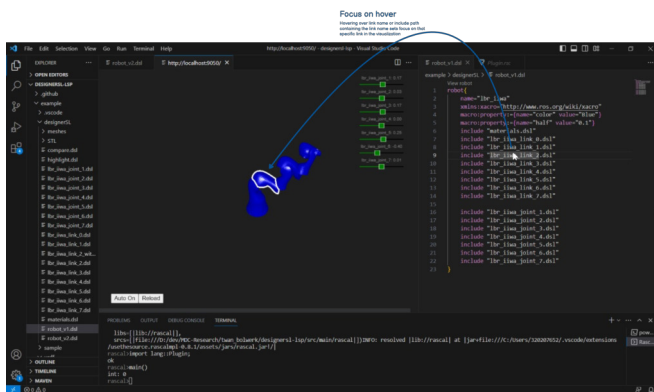


Fig. 5. Focus on hover link

⁹<https://www.blender.org/>

¹⁰<https://github.com/javahacks/vscode-urdf-viewer>

¹¹<https://www.babylonjs.com/>

C. Show differences between robot versions

The next use case is showing differences between two robot versions. When the GSL instances are defined as the corresponding physical links of the robot are displayed as solid while the remaining links become transparent, offering a visual distinction.

Listing 6 provides a GSL instance to compare two robot versions. In Figure 6, we show the same GSL instance on the right and on the link we have the graphical view with differences.

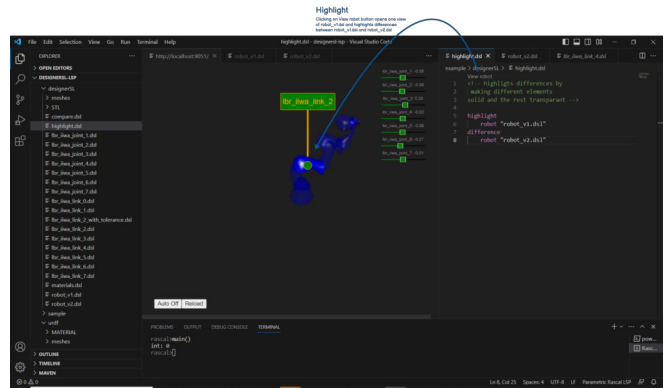


Fig. 6. Highlight robot

D. Present robot

The first use case we are going to describe is how we present the robot. The conversion tool has been executed and the VS Code URDF viewer is running. Salix is utilized to present the robot by a polling feature that performs an action over a specified interval as short as one second. Although this feature may impose certain performance overhead, it helps our language to fulfill the “liveness” quality criteria by consistently checking for differences and updating the view accordingly. A toggle button labeled “auto” is provided, as shown in Figure 5, which allows users to enable or disable the polling feature. This functionality offers a option to conserve system resources. With polling enabled the automatic updates are enabled, thus allowing for immediate feedback.

Upon detection of any discrepancy, a message with the updated model is send to our viewer which runs on a separate webserver thread. Rather than refreshing the complete web-view, the web-server can update the visualization gradually. This ensures that the robot is always maintained in the view during these gradual updates, reducing distractions from changing visual context. This feature enhances the user experience by providing a seamless transition during updates and maintaining continuity in the visualization.

Each time the user changes the source code and then saves the source code, a Rascal “summarize” event is triggered. In this event we set a flag to true, the polling mechanism that runs on a separate web-server thread then sends a HTTP message to our viewer. The viewer is equipped with a so-called listener (hooks) which update the visualization while preserving state.

Component	Responsibility
ViewerJS and UrdfSL	Visualization of robots and changes
Babylon JS	3D visualization
SalixJS and Salix	Bi-directional navigation
TypePal	Checking path existence and navigation
LanguageServer	Integration with the IDE, using events (summarize, document, lenses)
IDEServices	Opening files in the IDE or opening interactive content
XacroSL	Expression evaluation and resolving include path, using Xacro
GeometrySL	Providing visual semantics and URDF conversion

TABLE I
RESPONSIBILITY PER COMPONENT

This significantly reduces distraction of reloading graphics and therefore improving usability.

The “polling” mechanism is used to focus on an element that is hovered over by the mouse. This improves the user’s understanding of where the user is in source code. The focus mechanism sets the camera focus on that specific element and marks the element with a specific outline color as shown in Figure 5.

In this approach, a custom Xacro parser has been developed to convert Xacro code into GeometrySL. The shell “exec” function from Rascal is leveraged to call the Xacro compiler.

E. Tool’s components diagram

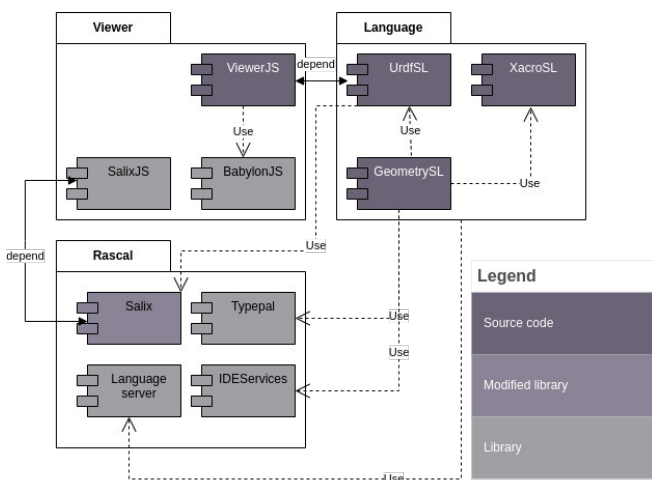


Fig. 7. Components

The component diagram in Figure 7 illustrates the software components and their relationships.

Each component shown in Figure 7 is mapped to its respective responsibility in Table I.

F. Link origin tracing

In our work, we trace link origin throughout the transformation process from GeometrySL to Xacro and finally to URDF, and vice versa. This approach draws inspiration from Inostroza et al. [15], where they track string origins during

transformations. However, our adaptation involves storing origin locations from physical links of the robot instead of strings and preserving this information across three transformations instead of one. This feature enables seamless navigation between the URDF source, GSL instance, and graphical representation of the robot model. As demonstrated previously, hovering over textual elements triggers the graphical element to be highlighted with a white outline. When holding down the Ctrl-key and clicking on a visual element, the user is directed to the corresponding section in the URDF source code. Conversely, clicking on a visual element without holding down Ctrl-key directs the user to the corresponding section in the GSL instance. With this approach bi-directional navigation is created. Upon hovering over text, the camera dynamically adjusts to center the active link, enhancing visual focus. When clicking on the link, a text editor is triggered, displaying the corresponding element’s description for detailed examination. The tracing of robot link elements is chosen because links contain the visually represented graphical mesh, allowing for bi-directional navigation with the graphical representation.

By re-using the shared abstract data structure of Xacro, the compatibility with Xacro is maintained and the development time of GeometrySL is reduced. The shared data structure, illustrated in Listing 7, encapsulates both Xacro and GeometrySL and is located in the Shared folder of the class diagram.

```
data Id_ = id(str id);
```

```
data XACRO_Attribute = attribute(Id_ \type, str val)
  | xacro_attribute (Id_ \type1, Id_ \type2, str val);
```

```
data XACRO_Object = ... // irrelevant alternatives omitted
  | link (Id_ name,
    list [XACRO_Attribute] attributes,
    list [XACRO_Object] elements,
    loc origin = unknown:// //); // link origin tracing
```

```
data XACRO = robot_(list[XACRO_Attribute] attributes,
  list [XACRO_Object] elements,
  loc origin = unknown:// //); // link origin tracing
```

Listing 7. Shared XACRO datastructure

An important point to highlight in Listing 7 is that all objects and attributes are generic, allowing for easy extension of new robot property semantics such as tolerances. However, in order to support bi-directional navigation a more concrete data structure is needed. The link is specifically specified to incorporate storing origin locations.

```
tuple[XACRO_Object,
  map[str,XACRO]] translate(Object obj,
  map[str,XACRO] includes){
  <xacro_obj,includes> = translate(obj.\type, obj, includes);
  if (xacro_obj.\type.id == "link") {
    name = get(1. attributes, "name");
    return <link (id(name),
      xacro_obj. attributes,
      xacro_obj.elements,
      origin=obj@loc), // keep track of link location
    includes>;
  }
  return <xacro_obj, includes>;
```

}

Listing 8. GeometrySL::Semantics

By opting for this generic data structure, it accommodates all valid XML formats, enabling the parsing and semantic translation of custom properties. However, a trade-off of this approach is that specific tasks like determining link origins requires iterating through all, as illustrated in Listing 10. Furthermore, the semantics of the generic data structure do not verify the validity of elements, accepting all inputs, whereas the UrdfSL semantics, as depicted in the Appendix' Listing 13, are more specific and restrictive.

```
map[str,loc] linkLocationMap = ();

map[str, loc] gatherLinkLocation(map[str,XACRO] xacros)
{
  map[str, loc] result = ();
  for (key <- xacros){
    result += gatherLinkLocation(xacros[key]);
  }
  return result;
}

map[str,loc] gatherLinkLocation(XACRO xacro){
  map[str, loc] result = ();
  for (obj <- xacro.elements){
    // pattern match on link elements
    if (link(_, _, _) := obj) {
      // map link name and location
      result += (obj.name.id:obj. origin);
    }
  }
  return result;
}
```

Listing 9. Gather link location algorithm

In Listing 9 the link location map is a relation that can be used in both directions. We can search on the link name but also look for its location to get the name of the link, a functionality that proves notably convenient.

G. Focus on hover

Different Integrated Development Environment (IDE) events are utilized, specifically employing the documenter event. This event has information about where the cursor of the user is located. This active cursor information is used to determine what link is being inspected. In our implementation we even were able to look up the active link through include statements.

```
bool isCursorLocationInLocation(loc cursor, loc linkLocation){
  return cursor.begin >= linkLocation.begin &&
  cursor.end <= linkLocation.end;
}
```

Listing 10. Cursor location algorithm

In the link translation from GeometrySL to URDF we keep track of the identifier, the name of the link and its location see Listing 10. This goes in two directions, hence bi-directional. The identifier is used to find the corresponding location in the link location mapping and in the other direction to find the

identifier based on its location. This location lookup checks if a certain location is in the same file and in between the row and column. If this is the case, it will return its identifier.

H. Highlight differences

The functionality of highlighting differences compared to previous version(s). Link origin tracing must be ignored, in order to strictly check for value equality, this is different opposed to [31] where origin is actually added and used to ensure file equality. The links are hashed such that they can be compared and stored more efficiently, see *removeOriginFromLink* function Listing 11.

```
map[str,str] removeOriginFromLink(list[URDF] links){
  map[str,str] result = ();
  for (link <- links){
    str uniqueHash = md5Hash(link.attributes + link.elements);
    result += (getName(link).val: uniqueHash);
  }
  return result;
}
```

Listing 11. Remove origin from link implementation

The URDF data structure consists of concrete data types for each URDF property. Each property has elements and attributes, parsed from the URDF robots. In order to compare robot1 (r1) and robot2 (r2), we extract from the elements mapping the “link” and store these in l1 and l2 respectively, such that we compare links only. Recall the explanation of the build-in `map` functions in Section IV-A. With `rangeR` we exclude the links in robot1 (r1) that do not exist in the robot2 (r2). Next the domain function is applied to the result, such that only the link identifiers are returned, since the keys are the link names, see Listing 11.

```
list [str] compare(URDF r1, URDF r2){
  l1 = getAll(r1, "link");
  l2 = getAll(r2, "link");
  return toList(domain(rangeR(
    removeOriginFromLink(l1),
    range(removeOriginFromLink(l2))));
}
```

Listing 12. Compare algorithm

Note that we chose to apply the algorithm to the URDF data structure instead of GeometrySL or XacroSL due to the URDF's single-file nature, simplifying the process. We check for changes with the algorithm in Listing 12.

VI. RESULTS

To see the final product in action, it is best to watch the demonstration videos. Due to confidentiality, we use the Panda robot instead of the Philips IGT interventional X-ray system. Figure 8 about bi-directional navigation¹² & figure 9 hover and highlight differences¹³ are videos of our tool in use.

¹²https://www.youtube.com/watch?v=n71kg1OKVus&ab_channel=fedesis3391

¹³https://www.youtube.com/watch?v=o_bJ8NsEODQ&ab_channel=fedesis3391




Fig. 8. video of the GeometrySL tool bi-directional navigation

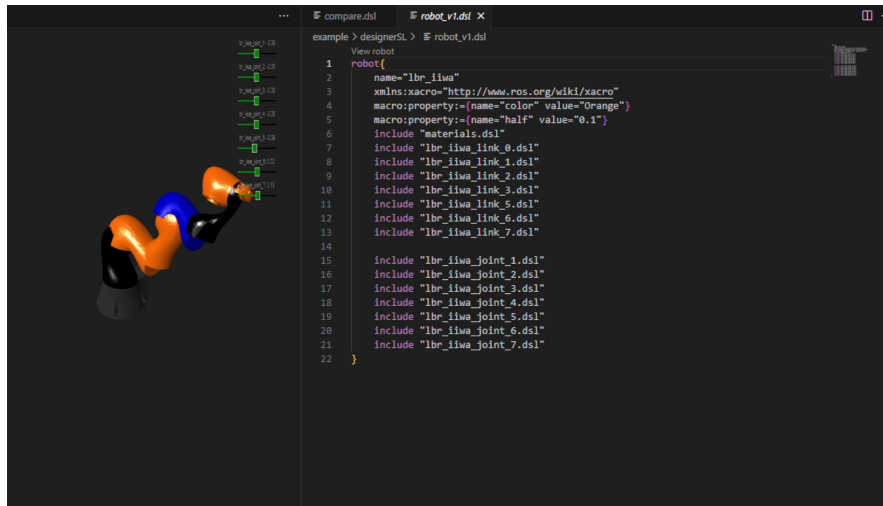


Fig. 9. video of the GeometrySL tool hover and highlight

VII. DISCUSSION

In contrast to the existing literature reviewed in Section II, which predominantly concentrates on visually representing properties of robots which can be visually represented such as positions of components and movements, our research places a strong emphasis on highlighting changes that cannot be visually represented, such as inertia and mass.

Leveraging the widely adopted URDF standard tapped into the existing knowledge base of end-users, simplifying their adoption and adaptation to the DSL. This underscores the notion that harnessing established standards can expedite the learning curve for new DSLs. Furthermore, through the incorporation of Xacro, we achieved composability within the language. Moreover, the development of GeometrySL expanded its capabilities with a visual language that empowers hardware engineers to precisely define how the robot is visually represented. This extension encompasses all known keywords and components while eliminating the cumbersome XML syntax, enhancing the user experience.

A. Bi-directional navigation

The benefits of bi-directional interaction, as outlined by Witte et al. [35], is that it reduces cognitive load and improves the overall user experience. By integrating the view and text editor through a bi-directional approach, our tool aims to alleviate cognitive load and further enhance the user experience. It allow users to interact with visual links by clicking on them, instantly directing them to the specific location in the text editor where that element is specified.

Salix has been used in game development to improve collaboration among multiple disciplines [33], resulting in improved collaboration. By leveraging the Salix library, a successful prototype of bi-directional input visualization was achieved. This feature enables users to effortlessly click on visual elements, which then redirects them to the corresponding location within the text editor.

Building upon the insights from [14], our approach embraces the concept of richness by not only providing visual feedback but also integrating the textual and graphical repre-

sentations. Although persistent changes are not supported via the visualization interface, this integration fosters a cohesive and intuitive user experience by establishing a strong connection between the two interfaces.

B. Immediate feedback

Earlier research [9] and [13] have indicated that immediate feedback significantly improves debugging. More recent studies also successfully applied immediate feedback in their DSL [32], [20], [3], [18]. Additionally, incorporating immediate feedback ensures that the DSL adheres to the liveness quality, improving the programming experience.

In VS Code, the “summarize event” of the Rascal Language server library is triggered each time the file is saved. We take advantage of this event to re-render the visualization using the latest valid state of the DSL, ensuring an up-to-date representation of the data.

This approach of fully refreshing the visualization, causes a brief moment of disappearance and reappearance of the robot, upon saving the file. This can be improved in terms of user experience. A more gradual change that maintains the robot in the exact same state and keeps it constantly in view would provide a smoother and more user-friendly way of communicating the modifications, minimizing context-switches and enhancing overall usability. By avoiding the robot’s disappearance, users can maintain continuous visual feedback and better understand the impact of their changes.

It makes sense to make this part of the summarize behavior the visualization serves a similar purpose as source code errors and warnings, and most languages perform these checks upon saving to achieve it. Additionally, this approach helps minimize the performance impact of running resource-intensive processes with each change.

Our language exhibits key qualities such as liveness, richness, and composability, as discussed in Horowitz et al.’s work on live programming [14].

Like mentioned in Munzner’s book [19], we aim to create a solution that addresses the invisibility of property changes and enhances user experience by implementing these highlighting strategies while considering potential conflicts with user-defined materials, overlapping color use, or making components transparent for highlighting. We opted against animations due to their potential to cause change blindness, distracting users from subtle property changes. Instead, we focus on highlighting the changed links to clearly differentiate changes, enhancing user experience by ensuring modifications are easily perceived and understood within the graphical representation. The highlighting differences aim to make changes more distinguishable and improve the user experience.

TMDiff [31] utilizes an algorithm that relies on source location to detect changes, similar to existing tools like Linux’s built-in diff and git diff. These tools typically compare lines of code, analyze differences on a line-by-line basis, and categorize changes as modified, deleted, or added.

In contrast, our solution specifically addresses the unique characteristics of URDF and GLS files, which represent

robotic systems with distinct links and joints identified by their names. We introduce a novel approach by hashing the links and joints including their unique identifiers while excluding their source location into a hashmap. This allows us to efficiently determine whether links or joints already exist in the system, enabling effective management of modifications and comparisons.

Moreover, our visualization tool offers enhanced capabilities compared to TMDiff. While both systems can highlight changes per link and between linked entities (in the case of joints), our solution goes further by seamlessly accommodating rearrangements of links and joints. This means that even if the structure of the system is altered, our tool can accurately compare the old and modified versions, providing a more robust and flexible comparison mechanism for this use case.

In conclusion, our study has demonstrated that integrating liveness, richness, and composability into a DSL tailored for hardware engineers at Philips IGT, featuring enhancements like bi-directional navigation, live graphical feedback, and the inclusion of robot components, can significantly enhance usability and effectiveness. Furthermore, it has underscored the challenge of visualizing invisible hardware properties, which often hinges on personal preferences and perceptions. These findings make meaningful contributions to the ongoing development of DSLs within this domain, emphasizing the critical role of user feedback in designing intuitive graphical feedback languages.

VIII. CONCLUDING REMARKS

We improved the hardware development workflow at Philips IGT by creating a textual Domain Specific Language (DSL) called GeometrySL or GSL. This DSL was used to formalize handovers from mechanical engineers to mechatronic engineers, preventing mistakes during the exchange process.

The novel approach of leveraging industry-standards URDF and Xacro, alongside techniques such as origin tracing and the LWB Rascal has resulted in a live graphical feedback on differences between Cyber-Physical System (CPS) versions. It enables bi-directional navigation among the graphical representation, URDF and the GSL itself, enhancing the efficiency and usability of the language.

The development of GeometrySL serves as a valuable case study that demonstrates the practical application of immediate, bi-directional visual feedback within an engineering context. The lessons learned from this project can inspire future research efforts and innovations in domain specific languages with live graphical feedback for CPS.

This research mainly focuses on graphically representing textual differences of 3D robot models. In our evaluation of the tool we had diverse feedback. The main challenge we faced was how to visualize changes made to invisible properties such as tolerances, mass inertia. A future work idea is to answer this question and improve our research. Explore innovative methods to visualize changes made to non-visual properties such as tolerances, mass, and inertia in robotic models.

While transparency has been utilized, investigate alternative approaches that effectively convey these modifications without significantly altering the overall view of the robot.

APPENDIX

```

data URDF =
  robot(map[str, URDFValue] attributes,
        map[str, list [URDF]] elements)
...
  map[str, list [URDF]] elements) // joint
| axis(map[str, URDFValue] attributes,
      map[str, list [URDF]] elements) // joint
| transmission(map[str, URDFValue] attributes,
              map[str, list [URDF]] elements) // robot
| actuator(map[str, URDFValue] attributes,
          map[str, list [URDF]] elements) // transmission
| plugin(map[str, URDFValue] attributes,
        map[str, list [URDF]] elements) // robot, link, or joint
| counterbalance(map[str, URDFValue] attributes,
                map[str, list [URDF]] elements) // joint
| tolerance(map[str, URDFValue] attributes,
           map[str, list [URDF]] elements) // robot custom property
;

```

Listing 13. Fragment of UrdfSL Abstract Data Structure

REFERENCES

- [1] L. Addazi, F. Ciccozzi, P. Langer, and E. Posse, "Towards seamless hybrid graphical-textual modelling for UML and profiles," in *Modelling Foundations and Applications*, ser. Lecture Notes in Computer Science, A. Anjorin and H. Espinoza, Eds. Springer International Publishing, 2017, pp. 20–33.
- [2] N. Albergio, V. Rathi, and J.-P. Ore, "Understanding xacro misunderstandings," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6247–6252.
- [3] D. Alique and M. Linares, "The importance of rapid and meaningful feedback on computer-aided graphic expression learning," vol. 27, pp. 54–60, 04 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1749772818300435>
- [4] R. Alur, *Principles of cyber-physical systems*. MIT press, 2015.
- [5] M. Ben-Ari, *Principles of the Spin model checker*. Springer Science & Business Media, 2008.
- [6] T. Bolwerk, "Improving the workflow for hardware engineers at philips with a domain-specific language and graphical feedback," 2023. [Online]. Available: https://www.cs.ru.nl/masters-theses/2023/T_Bolwerk___Improving_the_workflow_for_hardware_engineers_at_Philips_with_a_domain-specific_language_and_graphical_feedback.pdf
- [7] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, "Extensible markup language (xml) 1.0," 1998.
- [8] W. Cazzola and L. Favalli, "Scrambled features for breakfast: Concepts of agile language development," *Communications of the ACM*, vol. 66, no. 11, pp. 50–60, 2023.
- [9] C. Cook, M. Burnett, and D. Boom, "A bug's eye view of immediate visual feedback in direct-manipulation programming systems," in *Papers presented at the seventh workshop on Empirical studies of programmers*, ser. ESP '97. Association for Computing Machinery, 10 1997, pp. 20–41. [Online]. Available: <https://doi.org/10.1145/266399.266403>
- [10] J. Cooper and D. Kolovos, "Engineering hybrid graphical-textual languages with sirius and xttext: Requirements and challenges," in *2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*, 09 2019, pp. 322–325.
- [11] M. Fowler, *Domain-specific languages*. Pearson Education, 2010.
- [12] J. Gray, S. Neema, J.-P. Tolvanen, A. S. Gokhale, S. Kelly, and J. Sprinkle, "Domain-specific modeling." *Handbook of dynamic system modeling*, vol. 7, pp. 7–1, 2007.
- [13] T. R. G. Green and M. Petre, "Usability analysis of visual programming environments: A 'cognitive dimensions' framework," vol. 7, no. 2, pp. 131–174, 06 1996. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1045926X96900099>
- [14] J. Horowitz and J. Heer, *Live, Rich, and Composable: Qualities for Programming Beyond Static Text*, 03 2023.
- [15] P. Inostroza, T. van Der Storm, and S. Erdweg, "Tracing program transformations with string origins," in *International Conference on Theory and Practice of Model Transformations*. Springer, 2014, pp. 154–169.
- [16] P. Klint, T. Van der Storm, and J. Vinju, "RASCAL: A domain specific language for source code analysis and manipulation," ser. Proceedings of the 2009 Ninth IEEE International Working Conference on Source Code Analysis and Manipulation. IEEE, 2009, pp. 168–177.
- [17] L. Kunze, T. Roehm, and M. Beetz, "Towards semantic robot description languages," in *2011 IEEE International Conference on Robotics and Automation*, 05 2011, pp. 5589–5595, ISSN: 1050-4729.
- [18] A. Lozano, K. Mens, and A. Kellens, "Usage contracts: Offering immediate feedback on violations of structural source-code regularities," vol. 105, pp. 73–91, 07 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016764231500012X>
- [19] T. Munzner, *Visualization Analysis and Design*. CRC Press, 12 2014, google-Books-ID: NkYCWAAQBAJ.
- [20] NDC Conferences, "Real-time prototyping using visual programming languages - rui martins," 10 2018. [Online]. Available: <https://www.youtube.com/watch?v=cAiFJEcqwM4>
- [21] G. Onwubolu, *Mechatronics: principles and applications*. Elsevier, 2005.
- [22] F. Pérez Andrés, J. de Lara, and E. Guerra, "Domain specific languages with graphical and textual views," in *Applications of Graph Transformations with Industrial Relevance*, ser. Lecture Notes in Computer Science, A. Schürr, M. Nagl, and A. Zündorf, Eds. Springer, 2008, pp. 82–97.
- [23] M. Schuts, R. Aarssen, P. Tielemans, and J. Vinju, "Large-scale semi-automated migration of legacy c/c++ test code," *Software: Practice and Experience*, vol. 52, no. 7, pp. 1543–1580, 2022.
- [24] M. Schuts, M. Alonso, and J. Hooman, "Industrial experiences with the evolution of a dsl," in *Proceedings of the 18th ACM SIGPLAN International Workshop on Domain-Specific Modeling*, 2021, pp. 21–30.
- [25] G. B. Shelly and M. E. Vermaat, *Microsoft Office 2010: Introductory*. Course Technology Press, 2012.
- [26] L. Shen, X. Chen, R. Liu, H. Wang, and G. Ji, "Domain-specific language techniques for visual computing: A comprehensive study," vol. 28, no. 4, pp. 3113–3134, 06 2021. [Online]. Available: <https://doi.org/10.1007/s11831-020-09492-4>
- [27] R. H. Shih, *Parametric Modeling with Creo Parametric 2.0*. Sdc Publications, 2013.
- [28] T. van der Storm, "Semantics engineering with concrete syntax," in *Eelco Visser Commemorative Symposium (EVCS 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
- [29] van Rozen, "Live game programming with micro-machinations and rascal," 11 2014. [Online]. Available: <https://www.youtube.com/watch?v=YzsKaJEX4D4>
- [30] R. van Rozen and J. Dormans, "Adapting game mechanics with micro-machinations: International conference on the foundations of digital games," 03 2014, publisher: Society for the Advancement of the Science of Digital Games.
- [31] R. van Rozen and T. van der Storm, "Origin tracking + text differencing = textual model differencing: International conference on model transformation," pp. 18–33, 07 2015, place: New York Publisher: Springer.
- [32] R. van Rozen, "Cascade: A meta-language for change, cause and effect," 11 2022. [Online]. Available: <https://ir.cwi.nl/pub/32568>
- [33] R. van Rozen, Y. Reijne, C. Julia, and G. Samaritaki, "First-person realtime collaborative metaprogramming adventures," 12 2021. [Online]. Available: <https://ir.cwi.nl/pub/31301/>
- [34] R. van Rozen and T. van der Storm, "Toward live domain-specific languages: From text differencing to adapting models at run time," vol. 18, no. 1, pp. 195–212, 02 2019. [Online]. Available: <http://link.springer.com/10.1007/s10270-017-0608-7>
- [35] T. Witte and M. Tichy, "A hybrid editor for fast robot mission prototyping," in *2019 34th IEEE/ACM International Conference on Automated Software Engineering Workshop (ASEW)*, 11 2019, pp. 41–44, ISSN: 2151-0830.
- [36] D. Xue and Y. Chen, *System simulation techniques with MATLAB and Simulink*. John Wiley & Sons, 2013.

An AI-empowered energy-efficient portable NIRS solution for precision agriculture: A pilot study on a citrus fruit

Giulia Cisotto

0000-0002-9554-9367

University of Milano-Bicocca,

Department of Informatics, Systems and Communication,

Viale Sarca, 336, 20126 Milan, Italy

Email: giulia.cisotto@unimib.it

Dagmawi Deleegn Tegegn

0000-0002-5031-7589

SeleTech Engineering Srl

Via Volturmo, 37, 20861 Brugherio (Monza Brianza), Italy

Email: dag.tegegn@seletech.com

Alberto Zancanaro

0000-0002-5276-7030

University of Padova,

Department of Information Engineering,

Via Gradenigo, 6, 35136 Padova, Italy

Email: alberto.zancanaro.1@phd.unipd.it

Ivan Reguzzoni, Edoardo Lotti

SeleTech Engineering Srl

Via Volturmo, 37, 20861 Brugherio (Monza Brianza), Italy

Email: {i.reguzzoni, edoardo.lotti}@seletech.com

Sara L. Manzoni

0000-0002-6406-536X

University of Milano-Bicocca,

Department of Informatics, Systems and Communication,

Viale Sarca, 336, 20126 Milan, Italy

Email: sara.manzoni@unimib.it

Italo F. Zoppis

0000-0001-7312-7123

University of Milano-Bicocca,

Department of Informatics, Systems and Communication,

Viale Sarca, 336, 20126 Milan, Italy

Email: italo.zoppis@unimib.it

Abstract—Smart agriculture has seen impressive progresses in monitoring the quality of the crop and early detecting the onset of pathogens. However, this is typically achieved through smart, expensive, and energy-demanding robots and autonomous systems. We propose an AI-empowered portable low-cost short-wave near-infrared spectroscopy (sw-NIRS) solution that allows non-destructive measurements from plants and vegetables. In this pilot study, we specifically targeted an orange fruit and showed that it is possible to classify its different parts through sw-NIRS in the range 1350-2150 nm by using AI models, exceeding 97% accuracy. Also, we explored the minimum amount of energy needed to reach such high classification performance. In the future, we aim to extend this investigation to other targets (e.g., bean plants), to develop AI architectures to more accurately model the physiological conditions of the target, and to create a network of sw-NIRS sensors to simultaneously monitor a large-scale crop.

Index Terms—Near-infrared spectroscopy, machine learning, AI, chemometrics, energy efficient, green technology, smart agrifood, precision agriculture.

I. INTRODUCTION

SMART agri-food has recently seen tremendous developments thanks to new generation sensing, networking, and data analytics, i.e., ICT and artificial intelligence (AI), technologies.

This is allowing experts in the domain to quantitatively, continuously, and precisely monitor the conditions of the crops [1]. Near-infrared spectroscopy (NIRS) is one of the most popular techniques employed in the field, as it has already shown great potential in analyzing the quality and composition of foods [1], the maturity of fruits [2], [3] and crops [4], as well as the stress conditions of plants [5]. Another important advantage of NIRS is that this technology is available as portable devices (a very good review on the most recent handheld spectrometers can be found in [6]) with very fast scanning times (in the range of a few seconds) [7]. However, current portable NIRS devices suffer from some important limits, e.g., relatively large mass (over a few kilograms) [8], [9], relatively large sizes (e.g., a few tens of centimeters), no continuous acquisition modality, and a spectral resolution rarely below 2 nm. Here, we present a pilot study where a new handheld extremely lightweight but accurate NIRS spectrometer is used to acquire spectra from an orange fruit. With the complementary application of AI and machine learning (ML) modeling techniques [5], [10]–[12], we were able to classify different parts of the fruit with very high precision (classification accuracy over 97%), using very low energy.




Fig. 1. Experimental setup and some representative targets. (a) *Point 1* for *Part 1* (whole orange). (b) *Points 1* to *3* for *Part 3* (white layer) and *Part 4* (orange layer).

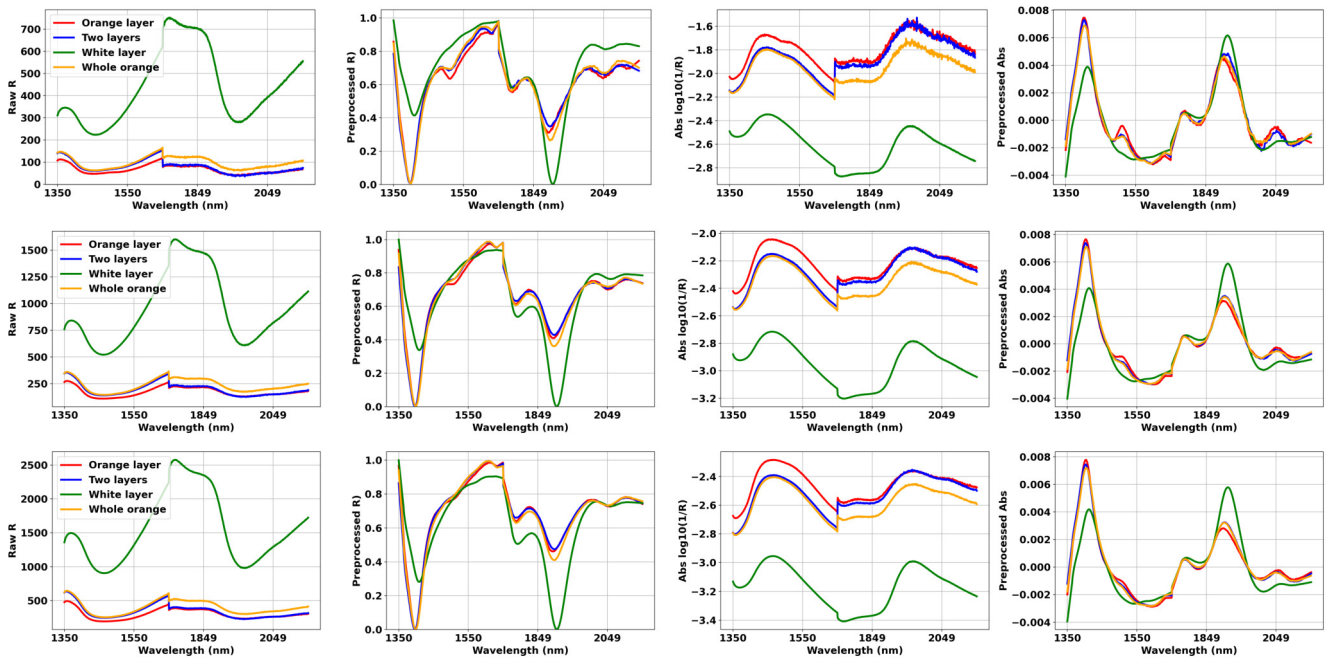


Fig. 2. Dataset in three single-lamp power settings at 100, 150, and 200.

II. MATERIALS AND METHODS

A. Portable NIRS

In this study, we employed a portable near-IR spectrometer that is able to capture the reflected light components from the target in a wide range of wavelengths, i.e., from 1350 nm to 2150 nm (technically defined as *short-wave NIRS*), with very high resolution of 1 nm. It is extremely lightweight (20 g) and small (45x25x13 mm).

Our spectrometer utilizes two microelectromechanical systems (MEMS) spectral sensors by Hamamatsu¹. These sensors

¹The datasheets of these sensors (C14272, C14273) can be found at <https://www.hamamatsu.com/eu/en/product/optical-sensors/spectrometers/mems-fpi-spectrum-sensor.html>

generate two independent spectra: one captures the spectrum from 1350 nm to 1650 nm, while the second captures from 1750 nm to 2150 nm. Although the sensors operate independently, the data from both of them are processed together in our analysis pipeline. This dual-sensor approach allows us to cover a wide spectral range (1350-2150 nm). Its power consumption is primarily determined by the lamp's power. An Os-hino lamp, integrated into the device, consumes approximately 60 mW at lamp power 100 and 270 mW at lamp power 250. Additionally, the SoC system and other operations consume 120 mW. This results in a total power consumption falling within the range of approximately 180 mW to 390 mW, which is exceptionally low for these handheld spectrometers [9].

This energy-efficient profile enables users to perform precise

spectral analysis without significant power demands, enhancing the spectrometer's practicality and versatility for a wide range of field and laboratory applications. The cost for such a hardware setup is very low (less than 500 \$ in the current device configuration), compared to many other available NIRS (e.g., hyperspectral) solutions (more than 1000 \$, depending on the specific features). Moreover, the device is prepared to embed a Bluetooth Low Energy (BLE) chipset for wireless transmissions (available in a later version). The acquisition time is very short, i.e., less than 2 s is enough to scan the entire range of wavelengths, making this device suitable for on-the-fly measurements of different types of targets, from plants to fruit, to animals [1].

B. Experimental protocol

Four different *parts* of an orange fruit were targeted for the NIRS measurement: (1) the whole orange fruit, (2) a two-layer target formed by the two outermost layers of the orange fruit, i.e., the orange and the white layers, (3) the white layer, and (4) the orange outermost layer. Spectra were collected from three different *points* on each part of the fruit. Fig. 1 shows the experimental setup and some representative targets.

Sixteen different lamp powers were available for the acquisition: from the value of 100 to the value of 250, with a 10-width step. Ten repeated measures were collected, in a continuous modality, on the same experimental condition, i.e., the combination of *part-point-lamp power*. The overall dataset finally included 1920 samples, with 702 features each (i.e., each feature representing one wavelength).

It is worth noting that we selected the three different *points* of every *part* with no rigorous localization criterion. This might have led to an increased *intra-part* variability but, at the same time, allowed us to prove that a ML model trained on these NIRS data is still able to distinguish across different parts, making our contribution closer to real-world applications.

C. Data preparation and pre-processing

From our previous investigations [1], we decided to use the standard normal variate (SNV) which consists in normalizing every spectrum by removing its own mean and dividing by its own standard deviation. This method aims to reduce the multiplicative effects of scattering and particle size, and allows to reduce the differences in the global intensities of the signals [13].

The Savitsky-Golay filter (SGF) is one of the most commonly used pre-processing steps in spectrometry and it consists of a 1-D filter that fits a polynomial function with degree p to a piece of data of length w . Often, first or second-order derivative is computed on the data before applying the filter. Based on our previous empirical investigations, we applied SGF on the reflectance values, setting $p = 2$, and $w = 30$ for MEMS1 while $w = 50$ for MEMS2 (to cope with the higher noise level), and using the first-order derivative. Thus, the first-order derivative emphasizes the dynamic changes in the reflectance spectrum, while the SGF smoothes up irrelevant

small peaks. Then, min-max normalization was applied to the filtered reflectance data to limit their values in the $[0, 1]$ range.

In this study, we investigated both the performance of our system when using all lamp powers (namely, *multi-lamp power setting*), or selecting one specific lamp power (namely, *single-lamp power setting*). In the latter case, during pre-processing, we extracted values obtained from the selected lamp power and performed the other operations on the reduced dataset, in the same way as in the case of the multi-lamp power setting. In both cases, we also investigated the possibility to reduce the dimensionality of the dataset (i.e., from the 702 dimensions) using principal component analysis (PCA). We decomposed the entire dataset in order to obtain a minimum number of principal components that explain the most variance in the dataset.

For visual inspection purposes, we also computed the absorbance spectra from the reflectance ones by computing $A = \log_{10}(1/R)$. The absorbance values were further filtered with a SGF with the same parameters values as the reflectance. All processing steps were implemented in Python 3.9, as available in the Jupyter environment in Google Colab.

D. Numerical experiments

Both for the *multi-lamp power setting* and the *single-lamp power setting*, we investigated the possibility of classifying the fruit parts, e.g., its different layers, from sw-NIRS data. Also, by visually inspecting the results of PCA, we observed the amount of separability among different classes while reducing the dataset dimensionality (originally set to 702) to the first 2 principal components.

Following the literature mainstream [4], we selected support vector machine (SVM) for the 4-class classification task, as it represents one of the most successful ML models in NIRS analysis. The Python class `sklearn.svm.SVC` was employed in the implementation of the classification task. This library is based on the `libsvm` package [14] that finds the best classification solution using a one-vs-one approach (default solver: C-SVC) in multi-class classification problems. During training, we applied a grid-search parameter optimization with the following values: kernel = [linear, radial basis function (rbf), polynomial], $C = [0.01, 1, 10, 100]$, $\gamma = [0.01, 0.1, 1, 10]$, where C is the regularization parameter that allows having a certain number of mis-classified samples, while γ represents the influence of far away samples in the computation of the separation hyperplane. We left the *degree* parameter for the polynomial kernel to the default value of 3. To train and validate the SVM model, we randomly selected 70% of the dataset, while keeping the remaining 30% for the test phase. During training, 5-fold cross-validation was applied.

We built the SVM models using pre-processed reflectance R values or, alternatively, using the first 2 principal components. However, in most cases, the models relying on the R values returned the best results. Thus, in the following, we show the classification performance obtained from the pre-processed R values, while the PCA results are used for visualization purposes, only.

To support the global effort of the scientific community in the direction of full reproducibility of research [15], [16], we share our Colab file at [17] and the dataset at [18].

III. RESULTS AND DISCUSSION

Visual inspection: Fig. 2 shows the dataset obtained by averaging 10 repeated measurements from three different points of every orange's part in three single-lamp power settings. It includes the raw reflectance spectrum, the pre-processed reflectance spectrum (i.e., as used for the subsequent analysis), the raw absorbance spectrum, and the pre-processed absorbance spectrum (as described in Section II-C).

We could observe that all spectra are consistent with previous literature [4], [8], [19]. Then, it can be noted that the white layer shows the most reflective intensity (both MEMS ranges), while the other three parts are much more similar to each other. Nevertheless, as one might expect, the whole orange and the two-layer targets produced spectra in between the white layer and the orange layer, with some slight differences in the two MEMS ranges.

Multi-lamp power setting study: When all measurements performed with any value of lamp power are included in the analysis, we found that the most explained variance is accounted for by the first 2 principal components, which represent 75% and 18% total variance in the dataset, respectively.

Furthermore, the best SVM classifier model (built on the pre-processed R values) was obtained with $C = 10$, $\gamma = 10$, and *rbf* kernel and reached 99.3% accuracy during the test.

Fig. 3 shows the normalized confusion matrix for the 4 classes in the multi-lamp power setting.




Fig. 3. Multi-lamp power setting: normalized confusion matrix for the 4 classes.

However, this setting leads to a high energy consumption that could reduce the battery life of the device. Thus, we performed a second stage of experiments while limiting the lamp power, i.e., the energy consumption of the device. We extensively tested the classification performance with every

single lamp power, in order to find the best trade-off between low energy usage and satisfactory classification accuracy values.






Fig. 4. Single-lamp power setting: normalized confusion matrix for the 4 classes for three different values of the lamp power (100, 150, 200, respectively).

Single-lamp power setting study: Fig. 4 shows the normalized confusion matrices for three different single-lamp powers, i.e., 100, 150, and 200. We can observe that, as the lamp power increases, the accuracy in the classification of the four orange's parts improves.

On the other hand, we also applied PCA to the three above configurations and noticed that for higher lamp powers, more variance is explained by the first principal components, only. Table I shows the explained variances for the first five principal components for the three different lamp powers.

TABLE I
EXPLAINED VARIANCE FOR THE FIRST 5 PRINCIPAL COMPONENTS FOR THREE DIFFERENT LAMP POWERS.

Lamp power	PC1	PC2	PC3	PC4	PC5
100	0.73	0.13	0.04	0.02	0.01
150	0.83	0.11	0.03	0.01	0
200	0.89	0.07	0.02	0.01	0

One might conclude that the lamp power acts as a relevant factor that increases the data variance. Including all measurements, regardless the lamp power, makes the variance spread more onto two principal components. However, the classification results are very satisfactory. On the other hand, it seems also possible that using a low-power lamp setting (i.e., 100) leads variance to be spread over more than one principal component. However, in this case, the classification performance decrease, possibly due to an increase of the variance that is not related to the target's characteristics.

Finally, Fig. 5 reports all classification accuracies, found during the test phase (over the pre-processed R values), for every lamp power value, i.e., from 100 to 250.

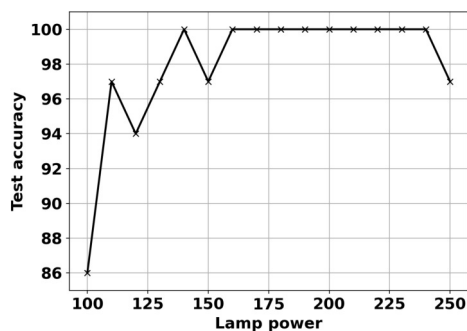


Fig. 5. Classification accuracy for every single-lamp power.

We could observe that, even with a single lamp power, it is possible to obtain very high classification accuracy values in the separation of the four orange fruit parts (chance level is 25%). Our portable spectrometer includes power-saving features when not in use, and switches to a higher power mode when acquiring spectra continuously. However, with this work we suggest that by empowering NIRS with AI (specifically, ML modeling), we could fasten the acquisition times and lower the energy consumption during the operating phase, while

obtaining very high classification performance. Specifically, from Fig. 5, one can choose to set the lamp power, for example 150, to obtain consistently good classification, consistently reduce the energy in the acquisition step and, at the same time, reduce the redundant variance not related to the target itself, as discussed above.

IV. CONCLUSIONS AND FUTURE PERSPECTIVES

The results presented in this study are still preliminary but made us prove the reliability and the advantages of an AI-empowered NIRS solution for the smart agri-food domain. In the future, we will extend this investigation to more complex scenarios: we will use the NIRS technology to monitor other target plants (e.g., bean and pothos plants), fruits, or to analyze organic compounds. Also, we can leverage the Bluetooth connection available on our spectrometer to deploy a network of wireless NIRS sensors [20] to precisely and promptly extract relevant information about the physiological conditions of plants in large-scale crops and new generation smart greenhouses. Furthermore, the current study was conducted in a semi-controlled lab environment as our aim was to perform a pilot study with our prototype. For real-world outdoor deployments, we will operate ad-hoc calibration based on the surrounding ambient light and we will perform a systematic feasibility study. Finally, we will the SVM model performance with other ML models and AI architectures (e.g., convolutional neural networks) [12]. Although there are other well-established solutions for precision agriculture in the market (e.g., RGB cameras), they often operate in different wavelength ranges, making direct comparisons with our system difficult. Therefore, future additional efforts will be dedicated to compare our system with more similar NIRS devices, currently not available in the market. Additionally, our system could be reviewed as a complimentary solution for satellite-based crop monitoring techniques [21] or drones-based weed mapping [22]. Then, to increase the sustainability of our system, we will investigate the actual amount of power consumption in the final NIRS product. However, with this first prototype, we were able to show the combination of hardware energy-saving features (e.g., battery size, lamp type, and other sensor configurations) with AI-based methods can significantly reduce the overall need of energy, while ensuring an effective monitoring of crops and plants for long periods of time.

ACKNOWLEDGMENT

This work was partially supported by the MUR under the grant "Dipartimenti di Eccellenza 2023-2027" of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Italy. AZ is also supported by PON 2014-2020 action IV.4 and GC by action IV.6 funded by the MUR.

REFERENCES

- [1] D. D. Tegegn, "Process of analyzing organic materials, based on processing of near-infrared spectra through advanced methods (PhD Thesis)," 2023.

- [2] P. Rodríguez, J. Villamizar, L. Londoño, T. Tran, and F. Davrieux, "Quantification of dry matter content in hass avocado by near-infrared spectroscopy (NIRS) scanning different fruit zones," *Plants*, vol. 12, no. 17, p. 3135, 2023.
- [3] K. Ncama, L. S. Magwaza, A. Mditshwa, and S. Z. Tesfay, "Application of visible to near-infrared spectroscopy for non-destructive assessment of quality parameters of fruit," *Infrared Spectroscopy-Principles, Advances, and Applications*, 2018.
- [4] A. M. Cavaco, D. Passos, R. M. Pires, M. D. Antunes, and R. Guerra, "Nondestructive assessment of citrus fruit quality and ripening by visible-near infrared reflectance spectroscopy," *Citrus-Research, Development and Biotechnology*, p. 95970, 2021.
- [5] A. Zancanaro, G. Cisotto, D. D. Tegegn, S. L. Manzoni, I. Reguzzoni, E. Lotti, and I. Zoppis, "Variational autoencoder for early stress detection in smart agriculture: A pilot study," in *2022 IEEE Workshop on Metrology for Agriculture and Forestry (MetroAgriFor)*. IEEE, 2022, pp. 126–130.
- [6] J. Müller-Maatsch and S. M. van Ruth, "Handheld devices for food authentication and their applications: A review," *Foods*, vol. 10, no. 12, p. 2901, 2021.
- [7] K. R. Borba, P. C. Spricigo, D. P. Aykas, M. C. Mitsuyuki, L. A. Colnago, and M. D. Ferreira, "Non-invasive quantification of vitamin C, citric acid, and sugar in Valência oranges using infrared spectroscopies," *Journal of Food Science and Technology*, vol. 58, pp. 731–738, 2021.
- [8] J. A. Cayuela and C. Weiland, "Intact orange quality prediction with two portable NIR spectrometers," *Postharvest Biology and Technology*, vol. 58, no. 2, pp. 113–120, 2010.
- [9] "Spectral Evolution NaturaSpec portable spectroradiometer," <https://spectralevolution.com/products/hardware/field-portable-spectroradiometers-for-remote-sensing/naturaspac-portable-spectroradiometer/>, accessed: 2023-09-25.
- [10] J. Martins, R. Guerra, R. Pires, M. Antunes, T. Panagopoulos, A. Brázio, A. Afonso, L. Silva, M. Lucas, and A. Cavaco, "SpectraNet-53: A deep residual learning architecture for predicting soluble solids content with VIS-NIR spectroscopy," *Computers and Electronics in Agriculture*, vol. 197, p. 106945, 2022.
- [11] Seletech Engineering Srl, "AI-empowered monitoring systems," <https://lnx.seletech.com/index.php/it/prodotti-e-progetti/intelligenza-artificiale/>, accessed: 2023-09-25.
- [12] W. Suphamitmongkol, G. Nie, R. Liu, S. Kasemsumran, and Y. Shi, "An alternative approach for the classification of orange varieties based on near infrared spectroscopy," *Computers and electronics in agriculture*, vol. 91, pp. 87–93, 2013.
- [13] J. A. Prananto, B. Minasny, and T. Weaver, "Near infrared (NIR) spectroscopy as a rapid and cost-effective method for nutrient analysis of plant leaf tissues," *Advances in agronomy*, vol. 164, pp. 1–49, 2020.
- [14] C.-C. Chang and C.-J. Lin, "LIBSVM: a library for support vector machines," *ACM transactions on intelligent systems and technology (TIST)*, vol. 2, no. 3, pp. 1–27, 2011.
- [15] F. Cabitza and A. Campagner, "The need to separate the wheat from the chaff in medical informatics: Introducing a comprehensive checklist for the (self)-assessment of medical ai studies," p. 104510, 2021.
- [16] K. Choudhary, D. Wines, K. Li, K. F. Garrity, V. Gupta, A. H. Romero, J. T. Krogel, K. Saritas, A. Fuhr, P. Ganesh *et al.*, "Jarvis-leaderboard: a large scale benchmark of materials design methods," *npj Computational Materials*, vol. 10, no. 1, p. 93, 2024.
- [17] G. Cisotto, "Python code associated with this publication at colab," https://colab.research.google.com/drive/14xU8u1Ao3smfnc_epTeIP9AS8etUjG8V?usp=sharing, last access: 2024-07-18.
- [18] G. Cisotto, D. D. Tegegn, I. Reguzzoni, and E. Lotti, "NIRS dataset associated with this publication," <https://github.com/CisottoGiulia/PON22-AI-NIRS-AgriFood>.
- [19] A. M. Cavaco, R. Pires, M. D. Antunes, T. Panagopoulos, A. Brázio, A. M. Afonso, L. Silva, M. R. Lucas, B. Cadeiras, S. P. Cruz *et al.*, "Validation of short wave near infrared calibration models for the quality and ripening of Newhall orange on tree across years and orchards," *Postharvest Biology and Technology*, vol. 141, pp. 86–97, 2018.
- [20] A. Zancanaro, G. Cisotto, and L. Badia, "Modeling value of information in remote sensing from correlated sources," *Computer Communications*, vol. 203, pp. 289–297, 2023.
- [21] P. Bertellini, G. D'Addese, G. Franchini, S. Parisi, C. Scribano, D. Zanirato, and M. Bertogna, "Binary classification of agricultural crops using sentinel satellite data and machine learning techniques," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 859–864.
- [22] G. Castellano, P. De Marinis, and G. Vessio, "Applying knowledge distillation to improve weed mapping with drones," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 393–400.

Automotive Cybersecurity Engineering with Modeling Support

Alexander Fischer
0009-0001-1737-7395
Nuremberg Institute of Technology,
Nuremberg, Bavaria, Germany
Email: a.fischer@th-nuernberg.de

Juha-Pekka Tolvanen
0000-0002-6409-5972
MetaCase,
Jyväskylä, Finland
Email: jpt@metacase.com

Ramin Tavakoli Kolagari
0000-0002-7470-3767
Nuremberg Institute of Technology,
Nuremberg, Bavaria, Germany
Email: ramin.tavakolikolagari@th-nuernberg.de

Abstract—Rapid advances of connected and autonomous vehicle technology have led to an increase in cyber-attacks. This in turn has driven the development of the ISO 21434 standard aimed at supporting the management of cybersecurity risks in the automotive industry. There is, however, a disconnect between the standard and the currently applied model-based development approaches that are increasingly applied for systems and software development. In this paper, we present tool support created for model-based automotive cybersecurity engineering. This tool is built upon the existing automotive systems development language, EAST-ADL, with extensions to address security in accordance with the ISO 21434 standard covering modeling support, calculation of security-related metrics such as impact, risk, and attack feasibility, and generation of ISO 21434 compliant security threat reports. Meeting the requirements of cybersecurity engineering according to ISO 21434 are demonstrated with two examples.

I. INTRODUCTION

THE DIGITIZATION and networking capabilities of modern vehicles require appropriate cybersecurity measures. As vehicles become more advanced, the risk of cyber-attacks increases, making it essential to identify and assess vulnerabilities in order to implement effective countermeasures. The ISO/SAE 21434:2021 standard “Road Vehicles—Cybersecurity Engineering” [1] provides appropriate means for identifying and assessing risk in the automotive industry, providing guidelines for identifying and assessing potential threats.

The importance of cybersecurity in the automotive sector is underscored by real-world incidents where attackers have exploited vulnerabilities to gain unauthorised access to vehicles (see Section II). These cases highlight the ease with which vehicles can be compromised due to insecure encryption systems or societal underestimation of the risk of an attack. In addition, the introduction of new connectivity features, such as infotainment systems, introduces additional attack vectors and presents new challenges in securing vehicle systems.

Model-based engineering is emphasised as the state-of-the-art approach in automotive software development. This methodology uses models to represent different aspects of the system, enabling the design, analysis and validation of complex systems. We describe the role of model-based approaches in more detail in Section III. Today, models of system and software development are typically kept separate from security models. Yet, integrating cybersecurity into the overall

system design is critical, especially with the increasing reliance on software components and the development of autonomous vehicle systems discussed in more detail in Section IV. Collaboration between system and security engineers is necessary to implement security-by-design principles. Models facilitate this collaboration by ensuring traceability of system functions and requirements, defining security objectives and analysing vulnerabilities.

The Security Abstraction Model (SAM, see Section VII) was developed with a focus on modelling cybersecurity threats and measures for automotive systems engineering. Recently, its scope has been extended to align with the ISO/SAE 21434 standard (see Section VI), which addresses cybersecurity in road vehicles. Although SAM offers a robust foundation for ISO/SAE 21434 by providing a metamodel that supports metric calculation and security threat reporting, there has been no tool support available for it until now.

The goal of this research is to show that tool support for SAM and model-based development is possible and can meet the requirements of ISO/SAE 21434 standard. We describe how tool support was developed as well as how it is applied. Our tool support is implemented in MetaEdit+ that enables collaborative development between systems and security engineering, see Section V. The developed tool features for cybersecurity engineering include in addition to modeling, calculation of security-relevant metrics such as impact, risk and attack feasibility, and the creation of ISO 21434-compliant reports. We demonstrate and show that the created modeling tool is viable and align with ISO 21434 with examples in Section VIII. In addition, many of the features described in ISO 21434 are also specified in other cybersecurity standards, so the security attack modeling components from SAM, including any that deal with social engineering attacks, are applicable to many relevant standards.

II. SECURITY RELEVANCE

The increase in cyber-attacks on vehicles reveals a need for action in the automotive sector to protect system components from external attacks. In particular, cases in which attackers were able to gain access to vehicles and start the engine by transmitting the radio key signal have been reported on by public media [2]. Vehicle owners must be made aware that

unlocking doors is easier than generally assumed and that even major manufacturers have been using insecure encryption systems for years [3]. In addition, the secrecy of data sheets does not ensure greater security, but makes thorough security verification more difficult [4]. The increasing connectivity of vehicles and the introduction of convenience features such as e.g. infotainment systems provide further attack vectors that lead to new challenges in identifying and addressing vulnerabilities [5].

Moreover, the trend towards vehicle-to-everything (V2X) communication, which includes vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) interactions, introduces additional risks. These systems rely on the exchange of information between vehicles and infrastructure to improve traffic flow and enhance safety. However, they also create new opportunities for attackers to intercept and manipulate data, potentially leading to collisions or traffic disruptions.

The proliferation of electric vehicles (EVs) also brings unique cybersecurity challenges. EVs often come with connected charging stations, which can be targeted to disrupt charging infrastructure or gain access to the vehicle's internal network. This not only affects the availability of charging but also raises concerns about the potential for large-scale attacks on the power grid.

In response to these threats, it is crucial for the automotive industry to adopt a multi-layered security approach. This includes implementing robust encryption methods, regular security updates, comprehensive testing of all systems, consideration of social engineering attacks, and an integration into the state-of-the-art automotive software systems development approach, i.e., model-based engineering, see the following section.

III. MODEL-BASED DEVELOPMENT IN THE AUTOMOTIVE DOMAIN

A key advantage of model-based development is its ability to enhance communication and collaboration among development teams. By using models, engineers from different disciplines (such as software, hardware, security, and systems engineering) can work together more effectively, sharing a common understanding of the system under development. This collaborative development approach helps to reduce errors and misunderstandings, leading to higher quality software and faster development cycles.

Model-based engineering currently represents the state of the art in the field of automotive software engineering. The primary reasons for model-based approaches are managing complex engineering tasks in better ways and effective communication [6]. In addition to support for collaboration it makes possible to design, analyze and validate complex systems by using models that represent different aspects of the system. It has proven to be extremely effective in supporting the development of automotive software as it enables the systematic design and analysis of functions [7]. The models are typically created with some general-purpose modeling

language like UML or SysML [8], [9], or with domain-specific languages targeting automotive systems like:

- Architecture Analysis & Design Language (AADL) [10] is a domain-specific language used for modeling the architecture of embedded systems, including automotive systems. It allows for the representation of both software and hardware components and their interactions. AADL is beneficial for performing performance analysis, such as timing and resource utilization critical in automotive applications.
- AUTOSAR (AUTomotive Open System ARchitecture) [11] is a standardized automotive software architecture framework that allows for the design and development of vehicle software with interoperability and scalability. It defines a set of specifications for software architecture, enabling the integration of components from multiple suppliers. AUTOSAR models are used to specify software components, their interfaces, and communication patterns, ensuring consistency and compatibility across different ECUs (Electronic Control Units).
- EAST-ADL (Electronics Architecture and Software Technology — Architecture Description Language) [12] is a domain-specific language tailored for automotive electrical and electronic systems. It provides a framework for modeling the architecture of vehicles, focusing on requirements engineering, functional analysis, dependability, and system design. It covers a more abstract design level compared to AUTOSAR. EAST-ADL supports the development process by linking requirements to design models and analysis tools, facilitating traceability and verification.

It is worth noticing that these well-known modeling languages applied in automotive do not recognize security, cybersecurity or support for ISO/SAE 21434:2021 standard (see languages used in the automotive industry [10], [11], [12]). We see that model-based development provides a basis for also supporting cybersecurity engineering and it can be done with an integrated manner. We introduce The Security Abstraction Model (SAM) in more detail in Section VII, as an extension to the EAST-ADL providing an integration and traceability between models of system development and cybersecurity.

IV. NEED FOR INTEGRATION OF SECURITY DESIGN INTO SYSTEMS MODELING

The increasing use of software components instead of mechanical components in vehicles and the development of autonomous vehicle systems require robust cybersecurity measures. Models allow system engineers and security engineers to collaborate and thus put the principle of “security-by-design” into practice. This collaborative modeling approach ensures that security considerations are embedded from the very beginning of the design process, rather than being retrofitted after the fact.

[13] gives an overview of the following advantages of the integrated approach:

- Models provide a structured way to document and trace system functions and requirements throughout the development lifecycle. By incorporating security requirements alongside functional requirements, engineers can ensure that security is treated as a core aspect of the system. This traceability allows better management of dependencies and the identification of potential security impacts arising from changes in system functionality.
- Security objectives need to be clearly defined to protect critical system assets and ensure the overall safety and privacy of vehicle occupants. Models can help in articulating these objectives in a precise manner, providing a clear roadmap for implementing necessary security measures. This includes specifying access control policies, data protection mechanisms, and secure communication protocols.
- To defend against potential cyber-attacks, specific security measures must be integrated into the system design. Models facilitate the systematic design and evaluation of these measures. For instance, threat modeling techniques can be used to identify potential attack vectors, and countermeasures can be designed and validated within the model. This proactive approach helps in mitigating risks before they materialize in the physical system.
- Continuous vulnerability analysis is crucial for maintaining the security of automotive systems. Models enable the simulation and analysis of various attack scenarios, helping engineers to understand the potential impact of different vulnerabilities. By analyzing these scenarios within the model, engineers can prioritize vulnerabilities based on their severity and likelihood, and implement appropriate mitigation strategies.
- The automotive industry is subject to stringent regulatory requirements regarding safety, e.g., ISO 26262, and security, e.g., ISO 21434. Integrating security within the design models ensures that the development process aligns with these regulations and industry standards. This alignment is essential for achieving certification and ensuring that vehicles meet legal and market requirements.

V. NEED FOR TOOL SUPPORT

Collaborative development work creating specifications, analyzing, checking and versioning them as well as transforming models to code, reports etc. requires tool support. In this paper we apply MetaEdit+ tool [14] to create and use modeling support for cybersecurity. MetaEdit+ is applied because it already supports existing automotive system development languages such as EAST-ADL and AUTOSAR. Second reason for using MetaEdit+ is that it can generate code directly from the models as well as allows creating generators for various purposes other than producing code, like checking, reporting, as well as producing input to other tools like simulators and analysis tools. This function not only provides considerable time and cost savings in development effort, but also improves the overall quality of the system developed.

Thirdly, and crucial for our work on security modeling, MetaEdit+ can extend and combine languages via metamodels, as well as create new domain-specific modeling languages. This flexibility allows for the customization of modeling languages to suit specific domain requirements. Once a metamodel is defined, developers can use it as their domain-specific language for modeling [15].

In Section VIII we describe how modeling support was created by defining security-related language concepts, rules and notation. We also present the generators that calculate security scores and produce relevant security documents as in ISO 21434. We demonstrate resulting tool support with examples.

VI. ISO STANDARD 21434

ISO/SAE 21434 contains objectives, requirements and guidelines related to cybersecurity engineering and can be used to implement a cybersecurity management system that also involves cybersecurity risk management [1]. The standard specifies the technical requirements for managing the cybersecurity risk of electrical and electronic systems (E/E-Systems) in road vehicles, including their components and interfaces. No specific technologies or solutions for cybersecurity are prescribed. ISO/SAE 21434 mandates risk treatment for all identified risks using classical options: risk avoidance, reduction, sharing, or retention and permits risk acceptance up to a defined threshold, as long as the decision is documented along with the retained risks [16]. According to ISO 21434, road vehicle cybersecurity is achieved when assets are adequately protected against threat scenarios. Assets worthy of protection include the various tangible and intangible components of systems such as software and hardware components, sensitive information and communication links. Threat scenarios are the potential cause for the compromised protection objectives of one or more assets [1]. ISO 21434 defines item as one or more components that implement a function at vehicle level, whereby a component is defined as a logically and technically separable part [1]. The item definition defines the target development system, which is subject to a cybersecurity-oriented development process, as precisely as possible and specifies the physical limits of the system under consideration as well as the areas to be protected. Based on the item definition, a threat analysis and risk assessment (TARA) is carried out from the perspective of affected road users. It serves to systematically identify threats and analyze the attack and defense mechanisms in the examined system and essentially consists of the following elements:

- 1) Item Definition [1, section 9.3]
- 2) Asset Identification [1, section 15.3]
- 3) Identification of Threat Scenarios [1, section 15.4]
- 4) Impact Rating [1, section 15.5]
- 5) Attack Path Analysis [1, section 15.6]
- 6) Attack Feasibility Rating [1, section 15.7]
- 7) Risk Value Determination [1, section 15.8]
- 8) Risk Treatment Decision [1, section 15.9]
- 9) Cyber Security Goals [1, section 9.4]

- 10) Cyber Security Claims [1, [RQ-09-06]]
- 11) Cyber Security Concept [1, section 9.5]

Cybersecurity engineering analysis identifies and explores potential actions that an abstract attacker could perform maliciously and the damage that could result from compromising the cybersecurity of a vehicle's E/E systems. Cybersecurity monitoring, remediation and incident response depend on changing environmental conditions, i.e. there is a constant need to identify vulnerabilities in road vehicle E/E systems and counteract new attack techniques.

The abbreviation CAL stands for Cybersecurity Assurance Level and, similar to the ASIL (Automotive Safety Integrity Level) in the ISO 26262 standard, is used to appropriately adjust the effort and care required for subsequent activities in the area of cybersecurity. The ISO/SAE 21434 standard specifies that an appropriate CAL should be defined for each threat scenario based on the associated impact and attack vectors. This is similar to setting risk values. While the risk value is dynamic and can change during the development process, the CAL is intended to remain stable during development as it is an integral part of a development requirement.

VII. SECURITY ABSTRACTION MODEL WITH EXTENSIONS

Security Abstraction Model (SAM) provides concepts for modeling security aspects of automotive systems. Figure 1 describes the metamodel of SAM illustrating which kind of security aspects are specified when modeling automotive systems with security considerations. In this figure, we present the complete metamodel so that the relationships between the entities become visible, as this is relevant for the reporting described later.

Originally SAM [17] did not recognize the later published ISO 21434 standard but this is now integrated into SAM and its metamodel [18]. This creates a link between the security requirements of the ISO 21434 standard and the models created based on SAM. Similarly, SAM was not originally developed explicitly for modeling social engineering but an extension has been developed that enables the modeling of social engineering attacks and maps the relationship of these attacks to the actors and the rest of the model [13]. These extensions enable a more comprehensive specification of cybersecurity aspects, their reporting as in ISO 21434 and calculating related metrics and scored. We describe these extensions in the next subsections, and their implementation to the modeling tool in Section VIII.

A. Integration with System Design

EAST-ADL [12] is a language for describing the system architectures of software-intensive automotive systems using an information model that represents technical information in a standardized way. The descriptions cover vehicle functions and features as well as functional and hardware architecture. The EAST-ADL model is structured according to abstraction levels, with each sub-model representing the relevant details of the complete embedded system of the respective abstraction level.

Security Abstraction Model and EAST-ADL are linked by the common concept of item. In EAST-ADL, item represents a functional or non-functional requirement of the system that is being described and modeled. SAM extends the concept of item by incorporating security properties. This enables SAM to specify security requirements that are necessary to fulfill the overall system requirements. These security requirements are integrated into the model to enable a comprehensive security analysis and to identify potential vulnerabilities and threats in the system.

Although SAM is developed as part of the EAST-ADL, it is not necessarily bound to EAST-ADL, offering flexibility in its application. SAM can be used independently of the rest of the system model to provide an overview of security-critical system parts before or at the beginning of the system engineering process. This independent utility allows engineers to identify and address potential security vulnerabilities early in the development cycle.

B. Scores

The latest version of SAM used at the time of writing includes a number of entities from ISO 21434 to enable a detailed risk assessment. These entities include Asset, Damage Scenario, Threat Scenario, ImpactRatingScore, RiskScore, AttackFeasibilityRating and AttackFeasibilityScore (see Figure 1). By integrating ISO 21434, not only can vulnerabilities now be assessed, but so can potential attacks and their impact on the system. For this purpose, the AttackFeasibilityScore, ImpactRatingScore and RiskScore are included in SAM. The AttackFeasibilityScore is calculated on the basis of the CVSS formula and makes it possible to estimate the feasibility of an attack [1].

In addition to the previously mentioned scores, ISO 21434 contains further scores that are available in the SAM metamodel. The ImpactRatingScore evaluates the impact of an attack scenario based on various factors such as the severity of the damage caused, the extent of the impact on the system and the potential duration of the impact. The RiskScore assesses the overall risk associated with a particular threat scenario. It takes into account the probability of a successful attack and the possible consequences of the attack. The RiskScore makes it possible to prioritize potential threats and take appropriate security measures to reduce or control the risk. The combination of these scores in SAM enables a more comprehensive security analysis and risk assessment. By taking into account vulnerabilities, attack scenarios, ImpactRatingScore and RiskScore, emerges a holistic picture of the security of a system in the automotive context. This makes it easier to identify potential risks and threats and take appropriate measures to increase the security and reliability of the system.

C. Social Engineering

SAM provides a basis for the assessment of social engineering attacks by including various scores and entities. A qualitative scoring system has been developed to specifically focused on social engineering. Integrating a scoring system




Fig. 1. SAM metamodel. View Online at <https://bitbucket.org/east-adl/sam/>

further improves the assessment and understanding of risks associated with social engineering, helping to develop appropriate security measures to minimize the impact of these attacks. Extensions to SAM were implemented at different levels, including new meta-entities, extensions to existing meta-entities, and supplementary documentation, enabling greater consideration of social engineering and standards.

Due to the existing integration of CVSS scores and other assessments in the Security Abstraction Model, it was necessary to investigate whether the extension and harmonization

would create redundancy. Redundancy is beneficial if there are reasons for mapping an issue in different ways. The social engineering entities were integrated into the metamodel to ensure a clear capture of security aspects without creating unnecessary duplication or repetition of metrics.

D. Reporting

Reporting is essential in the context of ISO 21434, which explicitly requires it. Specifically, a cybersecurity assessment report (RQ-06-31) serves as appraisal of the level of cy-

bersecurity. Although the standard does not provide explicit guidelines on the format or structure of such reports, our implementation of the report generation adheres closely to the principles described in ISO 21434. This ensures that the cybersecurity assessment report effectively communicates the findings and recommendations derived from the assessment process.

In accordance with ISO 21434, the report is primarily focused on assets, reflecting the standard's emphasis on asset-oriented cybersecurity management. Its structure is closely based on the example in Annex H of the standard. However, by integrating the social engineering aspect of SAM, we have also introduced reporting that focus on human actors and recognize the importance of the human element in cybersecurity. In addition, we have included a section dealing with miscellaneous items and how they relate to vulnerabilities and the associated vulnerability scores. While the report sections relating to social engineering and miscellaneous are not explicitly included in the standard, their inclusion broadens the scope of the report and provides stakeholders with a holistic reporting of cybersecurity risk and mitigation.

There are several advantages to automatic report generation:

- It enables the hierarchical organization of multiple models, facilitating the creation of comprehensive reports that cover different aspects. This hierarchical structuring enables a systematic and coherent presentation of information across different levels of abstraction.
- Automated report generation can incorporate item definitions by linking to EAST-ADL architecture models, providing insight into potentially at-risk vehicle features and their interrelationships. This integration increases the depth and specificity of the report.
- By selecting relevant properties, the calculation and reporting of scores is automatically generated, which ensures efficiency and accuracy. In cases where multiple values are applicable, these are aggregated, with the maximum value being reported.

VIII. TOOL SUPPORT: LANGUAGE DEFINITION AND USAGE

This section presents the tool support for security modeling. We first describe the implementation of support for SAM, including the modeling language, score calculators and the reporting of security threats in accordance with ISO standard 21434. Subsequently, we provide two examples demonstrating the use of the developed modeling tool, alongside score calculation, reporting and tracing to other system design models.

Our implementation of tool support began by extending the existing language definition of EAST-ADL and its associated security language. Although EAST-ADL is supported by various tools, we applied in MetaEdit+ the latest version of EAST-ADL (v2.2)¹. Since MetaEdit+ enables the co-evolution of metamodels and models [19], the changes made to the mod-

eling support were automatically updated to already existing models.

The language definition covered all parts needed for obtaining tool support: Not only the metamodel and related constraints, but also the notation, guidance for creating and editing models, as well as updating older versions or notifying modelers to make changes when automatic update were not considered feasibly, such as when there was a risk of losing relevant data. Finally, generators for various score calculations and threat security reporting were defined, in addition to those available in MetaEdit+ for EAST-ADL, such as Simulink, Hip-Hops and ReqIF, or defined by users targeting external tools like SPIN, UPPAAL, Stateflow and Reliability Workbench².

A. Metamodel Extensions

For modeling support, the metamodel of SAM was defined by two person with MetaEdit+ Workbench, and then tested by other modelers by using the same language with the modeling editors, browsers, and collaboration tools of MetaEdit+. We created several security models as test cases including the Brake-By-Wire example presented here later³. Figure 2 shows the elements of the security modeling language in MetaEdit+. The list of Objects shows the key modeling objects, the list of Relationships shows the connections between these elements, and the list of Roles shows how an object participates in the relationships, such as being directed or undirected, having constraints, or detailed properties.




Fig. 2. Extended SAM definition

Figure 3 details the definition of HumanActor, which has 11 properties. The first three are inherited from EAST-ADL

²<https://metacase.com/solution/east-adl.html>

³SAM implementation can be accessed at https://bitbucket.org/east-adl/sam/src/master/MetaEdit-Extension/Reporting_Examples

¹<https://east-adl.info/>

and AUTOSAR metamodels. These three properties have rules and constraints, such as 'Short name' being mandatory and starting with an alphabetical character followed by possible characters, numbers, or underscores and constraint with maximum length (defined as a regular expression: $[a-zA-Z][a-zA-Z0-9_]{0,127}$). These are followed by the characteristics of HumanActor in SAM (see Figure 1): Curiosity, Helpfulness, Credulity, Greed, Ignorance, and Carelessness – all of which are fixed value enumerations.

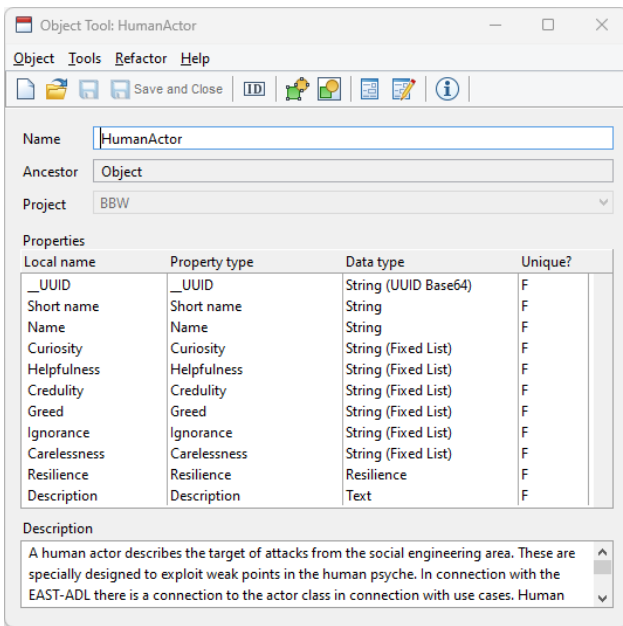


Fig. 3. Definition of HumanActor in metamodel of MetaEdit+

While the metamodel in Figure 1 identifies many language concepts as individual objects, such as resilience or metrics elements, we aimed to minimize the modeling effort not only in terms of creating model elements but also in terms of updating, deleting and checking specifications. As a result, the implementation as a modeling language exhibits some differences from the metamodel illustrated in Figure 1. This is mainly to minimize the modeling effort and improve usability. The main differences are:

- Since Resilience is a mandatory item for a human actor, it is a property of HumanActor. This way language user is expected to add – and later edit – just one element in a model rather than two and a connection between.
- The same approach is also applied for the 4 metrics elements: while they can be added to and visualized in the model, they are not mandatory. Modeling editor can calculate the metrics even if those metric elements are not explicitly added to the model. Figure 4 illustrates this in the user interface at the bottom of the screen by showing individual metric values for vulnerabilities and attacks yet showing CVSS basic and temporal scores for vulnerability as well as ISO 21434 feasibility score for

the attack directly in the diagram, which is what the user wanted in this case.

- Attack motivation is an element of the modeling language, and its subtype is selected from the property with mandatory value. Thus, the type of AttackMotivation (Harm, Financial Gain etc.) can be changed without deleting the old one and creating and re-connecting a new one.
- While the metamodel of SAM defines directed associations among security concepts, the modeling language does not expect models to be created in that order: the created editor shows the correct direction regardless of how the user opts to link model items. In other words, the model is created correctly independently of the order in which the modeler decides to create relationships.
- Default values for enumerations are provided.
- Properties of model elements are listed in the order that would be the order that would be most natural for considering the security properties.

We did not enforce all rules as mandatory, such as requiring each HumanActor to have a defined Resilience. Instead, we allowed for more flexibility in modeling, but we also provided guidance to language users to complete the security model. We defined 17 checks derived from the metamodel to provide warnings, which were shown to the language user during modeling. As an example, at the bottom of Figure 4 is shown a warning that SecurityConcept is not related to any Requirement. Additionally, we defined recommendations for creating security models that deal with optional links: linking Attacks to Actors, SecurityConcepts with Attacks, and DamageScenarios to Hazards – the last been shown as a recommendation by the tool for the security model in Figure 4.

B. Notation and Guidance

The security model example also illustrates the notation: How models are presented for humans to read, edit and use for communication. Our tool implementation therefore covered creating notation for the respective language elements. Figure 5 shows the definition of notational symbol for DamageScenario: It shows the name that user enters and impact rating score produced by the score calculator generators. The notation also shows the type of model element as a part of the notational symbol. Such guide is useful when creating or reading the models in the first place but for experienced modelers it becomes redundant text that consumes extra space and thus can be hidden by the language user from the diagrams if desired.

In addition to providing guidance during modeling, the defined metamodel thoroughly describes individual language elements. These descriptions are accessible directly in MetaEdit+ through the help system, which is available from the editor's toolbar or individually for each language element when in use.

C. Co-evolution

Given that SAM itself also evolved, we implemented guidance to update the existing security models to be compliant

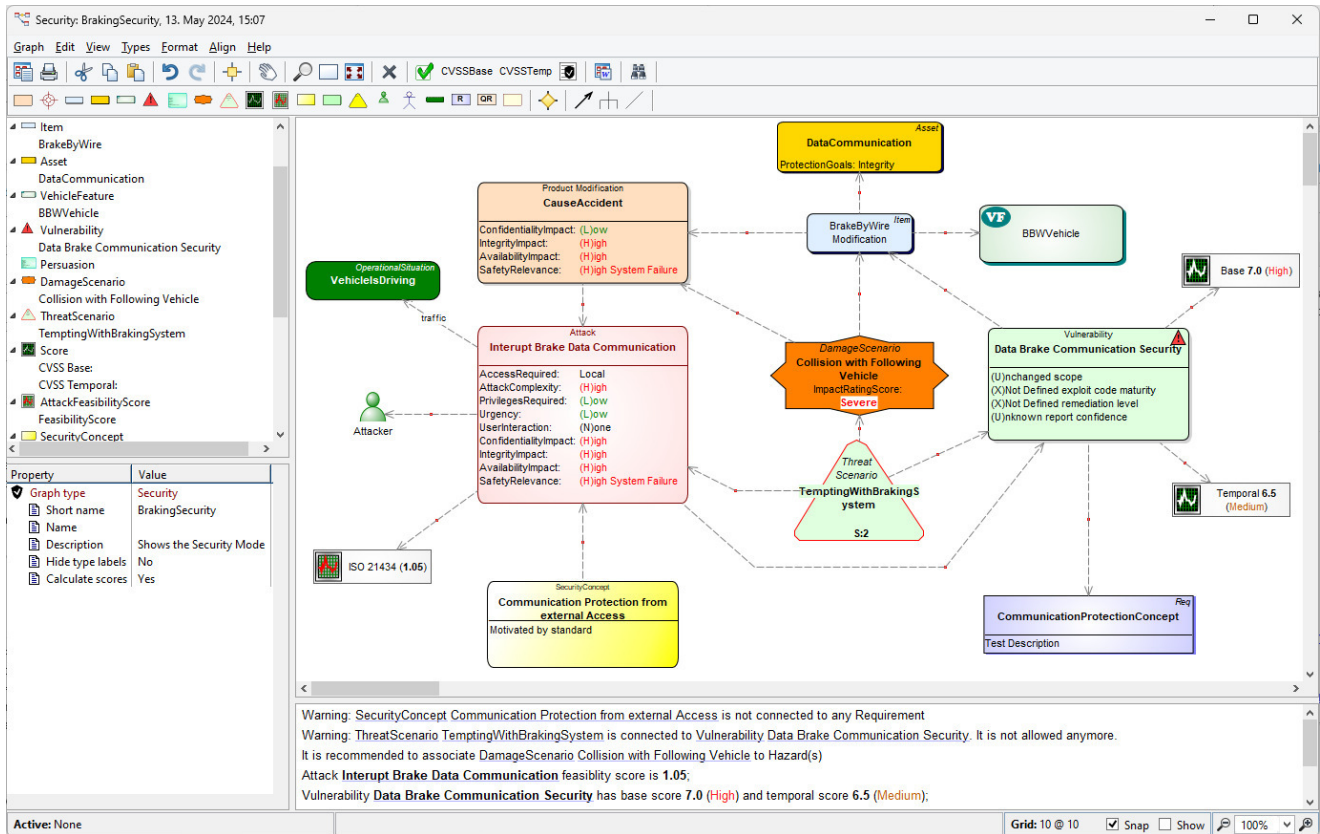


Fig. 4. Security model in modeling editor illustrating checks and recommendations

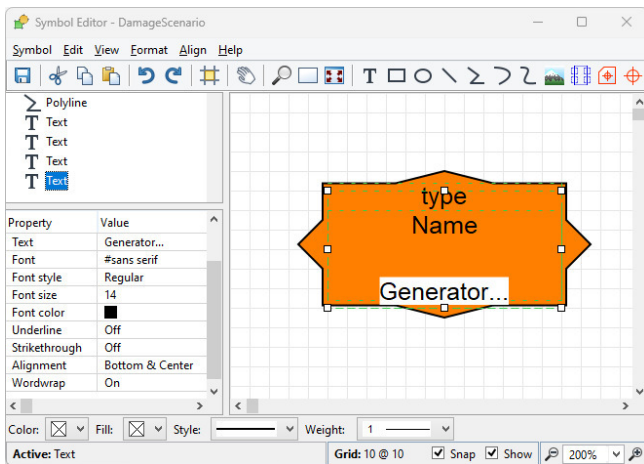


Fig. 5. Defining the notation for DamageScenario

with the latest language version. MetaEdit+ manages with in-built functionality automatically model changes that are caused by adding and renaming items in the metamodel, constraints and notation (for details see [19]). However, some changes to models may require human intervention when automatic updates are not feasible. To prevent the loss of critical model information, we followed a deprecation strategy:

existing models can still be used, but all new models follow the latest language version. Additionally, we implemented guidance within the modeling editor to assist users in updating their models. This feature is illustrated at the bottom of the editor in Figure 4: "Warning: ThreatScenario Tempting-WithBrakingSystem is connected to Vulnerability Data Brake Communication Security. It is not allowed anymore." Similar co-evolution support could be applied in the future when support for cybersecurity modeling evolves or new versions of the ISO standard or SAM are developed.

D. Metrics

As models provide rich details on security aspects, they can be used for various assessment purposes. We implemented support for SAM-based security models with the CVSS. Once a security model is created with the required data, the modeling tool calculates various scores automatically, like in Figure 4 for vulnerability of Data Brake Communication Security the score of Base CVSS is 7.0 (High) and Temporal CVSS is 6.5 (Medium) and for the specified attack ISO 21434 score is 1.05.

Since attacks can consist of subattacks, calculating vulnerability metrics must consider the whole attack subtree. In our implementation of CVSS, we considered the most severe case by recognizing the most severe attack within attack tree as a basis for calculation. The same principle is applied when

different types of individual attacks are related to the same vulnerability.

Scores on vulnerability and attack feasibility are calculated similarly at the time of modeling and illustrated either in the diagram or in a separate report pane below the diagram. Figure 4 illustrates scores at the bottom of screen and AttackFeasibilityScore next to the Attack element. The impact rating for Damage scenario (Severe) and risk score for ThreatScenarios (S:2) are also illustrated in Figure 4.

E. Documenting and Reporting

Existing documentation generators were available in MetaEdit+ for the purposes of reporting. These generators, however, did not recognise the needs of ISO 21434. Given that the SAM was made to recognize explicitly cybersecurity, we defined a threat reporting generator based on the reporting requirements (as in Section VII-D).

Figure 6 shows the result of this generator produced from Figure 4 and from the related system design specifying the vehicle features (Figure 7) and the system functions (Figure 8). Figure 7 shows a small part of the model specifying features related to the braking system. These features are realized by some design functions and hardware functions of EAST-ADL. Figure 8 illustrates a part of the logical design functions of the braking system that are also recognized in the generated security report. Both security report and metric calculators were implemented with generator system of MetaEdit+ [14].

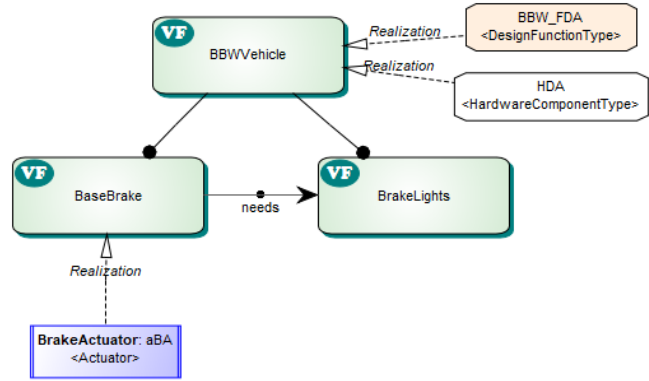


Fig. 7. Vehicle feature model: braking (fraction)




Fig. 8. Design level functions of braking system (fraction)

Standard ISO/SAE 21434
 There are currently 1 Items with 1 Features and 0 Human actors affected by overall 0 Hazards and 1 Attacks. 1 security models are not handled yet (BrakingSecurity). 2 of the Development Phases are affected. Item BrakeByWire is affected by the highest threat.

Item definition
 The item BrakeByWire: is linked to DesignFunctionType BBW_FDA (DesignFunctionArchitecture for Braking from Pedal to ABS) via is linked to HardwareComponentType HDA (Hardware architecture for braking system.) via feature BBWVeh

Asset identification
 Asset DataCommunication is related to DamageScenario Collision with Following Vehicle with the following se

Impact rating
 DamageScenario Collision with Following Vehicle has impact rating: Severe

Threat scenario identification
 DamageScenario Collision with Following Vehicle is related to ThreatScenario TempInqWithBrakingSystem (

Attack analysis
 ThreatScenario TempInqWithBrakingSystem is related to: Attack Interrupt Brake Data Communication (Test I

Attack feasibility rating
 Interrupt Brake Data Communication has attack feasibility rating 1.05

Risk value determination
 TempInqWithBrakingSystem has aggregated attack feasibility rating 1.05 and Impact rating Severe with risk

Targets
 Target BrakeByWire is related to Data Brake Communication Security

Vulnerability analysis
 Vulnerability Data Brake Communication Security has vulnerability base score 7.0 (High), and temporal score 6.5 (Medium)

Fig. 6. Sample of Security Analysis Report

Traceability from security models to system design is visible in the security analysis report. For example, in Figure 6 the item definition at the beginning of the report is linked to the design functions and hardware functions of the braking

system. Also the summary at the beginning of the generated report shows that security models are related to two different development phases of EAST-ADL, namely to the vehicle level in which features related to items are defined as well as to the design level functions realizing those features.

Figure 9 shows another report targeting analysis of social engineering threats in automotive systems. This report is generated from a security model shown in Figure 10 representing a social engineering attack that affects the braking system. It shows a baiting attack in which the braking system is compromised through deception maneuvers. The report identifies the human actors involved, their vulnerabilities and their resilience to such attacks. Additionally, it provides insights into the persuasion methods used in the social engineering attack, improving understanding of the potential dangers posed by human manipulation tactics. This holistic approach to reporting provides valuable insight into the intricacies of cybersecurity risks associated with social engineering and helps develop robust countermeasures to protect automotive systems from such vulnerabilities.

The reports illustrated in Figures 6 and 9 show that they provide links from reported items back to the security models and other system development models. This clear traceability shows that security aspects do not need to be addressed in isolation, but can be linked to the rest of the system development. These reports can be produced directly to external files like used for word processors or web browsers.

In addition to reporting on individual security model – as illustrated in the previous examples – security threat reporting is also available for all EAST-ADL models: It can be generated

Social engineering

There are currently 1 Human actors with 5 exploitable human weaknesses. 1 Principles of Persuasion are used in 1 attacks.

Human actors

Car Owner has the following properties:

Curiosity: (L)ow
 Helpfulness: (H)igh
 Credulity: (H)igh
 Greed: (N)one
 Ignorance: (L)ow
 Carelessness: (H)igh
 with resilience:
 Cautiousness: (H)igh
 Contentment: (L)ow
 Courage: (L)ow /intermediate
 Experience: (H)igh
 Knowledge: (H)igh

Persuasions

Related to Attack Interrupt Brake Data Communication through baiting

Persuasion Persuasion has the following properties:

Reciprocity: Obligation: (H)igh
 Concession: (H)igh
 Scarcity: (H)igh
 Authority: (H)igh
 Consistency and Commitment: ()intermediate
 Liking: (L)ow
 Social Proof: (L)ow

Miscellaneous

There are currently 1 Items and 1 Vulnerabilities. The highest vulnerability score is 6.2.

Fig. 9. Sample of social engineering report

for any selected hierarchy of EAST-ADL models combining multiple security models into a single security threat report. This capability enhances collaboration by allowing traces from system designs to be followed to all vulnerabilities and attacks across the entire developed system.

IX. CONCLUSION

This paper presents a tool support for model-based cybersecurity engineering in the automotive domain. It shows how tool support can meet the requirements of ISO/SAE 21434 standard in model building, calculating metrics and security threat reporting. Our tool, built on the EAST-ADL language with security extensions, provides a solution to support these model-based approaches. By integrating system and security modeling, along with capabilities for calculating security metrics and generating ISO-compliant reports, the tool enables engineers to navigate the complexities of automotive cybersecurity with confidence. Furthermore, the tool's ability to guide engineers in defining and integrating security models with system models underscores its user-centered design and practical utility.

The significance of this work extends beyond its immediate application in automotive cybersecurity. As the latest enhancements to the metamodel enable a complete representation of the ISO 21434 standard, it lays the groundwork for broader adoption across industries where cybersecurity standards are of highest importance. Moreover, the versatility of the exten-

sions, particularly those related to social engineering attacks, positions it as a valuable resource for compliance with various cybersecurity standards beyond ISO 21434.

While the modeling support is readily available our plan is to apply it to model various security cases to evaluate it and identify possible areas for extensions. Another direction for future research is to extend tool support, and possibly the metamodel of SAM, to support the latest versions of metric calculators like version 4.0 of CVSS.

Future research could investigate the use of Large Language Models (LLMs) to automatically generate models based on attack data. This approach has the potential to rationalize the modeling process and enable not only security engineers but also automotive engineers to contribute to the creation of security models. By automating the generation of parts of the models that currently require manual modeling, such as specific attack scenarios and vulnerabilities, significant time savings can be achieved.

Other extensions to the metamodel could relate to the implementation of specific mechanisms, such as cryptography. Although the metamodel already allows the modeling of requirements and security concepts, these additions could allow a more detailed and accurate modeling of the internal relationships of these mechanisms.

The Cybersecurity Assurance Levels (CALs) from the ISO 21434 standard can be specified in the tool for a security goal. However, these security goals and other entities from the concept phase, such as requirements, are not currently included in the reporting, as the current reports focus primarily on risk assessment. For CALs, it is important to note that no consensus has yet been reached on how to determine and treat such a parameter, so this aspect has been relegated to the Annex only [20]. This could be a potential future extension, allowing for the creation of reports that encompass requirements, security goals, and concepts, even though this is not explicitly required by the standard.

REFERENCES

- [1] "ISO/SAE 21434:2021, Road vehicles – Cybersecurity engineering." Aug. 2021. <https://www.iso.org/standard/70918.html>
- [2] J. Li, Y. Dong, S. Fang, H. Zhang, and D. Xu. 2020. "User Context Detection for Relay Attack Resistance in Passive Keyless Entry and Start System," *Sensors*, vol. 20, no. 16, p. 4446, Aug. 2020, doi: <https://doi.org/10.3390/s20164446>.
- [3] F. D. Garcia, D. Oswald, T. Kasper, and P. Pavlidès. 2016. "Lock it and still lose it – on the (In)Security of automotive remote keyless entry systems," in 25th USENIX Security Symposium (USENIX Security 16), ser. SEC'16. Austin, TX, USA, Aug. 2016, pp. 929-944.
- [4] L. Wouters, E. Marin, T. Ashur, B. Gierlichs, and B. Preneel. 2019. "Fast, furious and insecure: Passive keyless entry and start systems in modern supercars," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2019, no. 3, pp. 66-85, <https://doi.org/10.13154/tches.v2019.i3.66-85>
- [5] Costantino, A. La Marra, F. Martinelli, and I. Matteucci. 2018. "Candy: A social engineering attack to leak information from infotainment system," in 2018 IEEE 87th Vehicular Technology Conference (VTC Spring), pp. 1-5, <https://doi.org/10.1109/VTCSpring.2018.8417879>
- [6] H. Gustavsson, E. P. Enoiu and J. Carlson. 2022. "Model-Based System Engineering Adoption in the Vehicular Systems Domain," 2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS), Sofia, Bulgaria, pp. 907-911, <https://doi.org/10.15439/2022F47>.

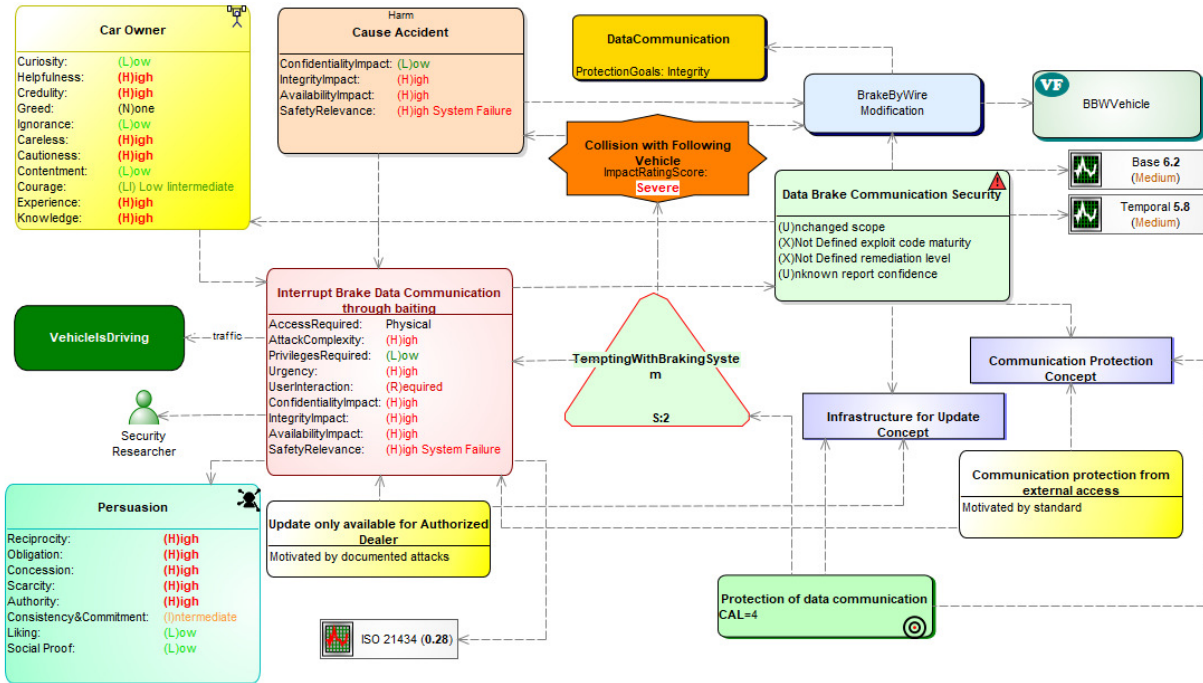


Fig. 10. Security model in modeling editor illustrating a social engineering attack

- [7] M. Broy, M. Feilkas, M. Herrmannsdoerfer, S. Merenda and D. Ratiu. 2010. "Seamless Model-Based Development: From Isolated Tools to Integrated Model Engineering Environments," in Proceedings of the IEEE, vol. 98, no. 4, pp. 526-545, <https://doi.org/10.1109/JPROC.2009.2037771>.
- [8] OMG. Unified modeling language specification version 2.5.1. 2017. <https://www.omg.org/spec/UML/2.5.1/>
- [9] OMG. Systems Modeling Language Specification version 1.6. 2019. <https://www.omg.org/spec/SysML/1.6/>
- [10] P. H. Feiler, D. P. Gluch and J. Hudak. 2006. "The Architecture Analysis & Design Language (AADL): An Introduction," <https://doi.org/10.1184/r1/6584909.v1>
- [11] AUTOSAR: Enabling Continuous Innovations. 2024. <https://www.autosar.org/>
- [12] H. Blom, H. Lönn, F. Hagl, Y. Papadopoulos, M.-O. Reiser, C.-J. Sjöstedt, D.-J. Chen, F. Tagliabò, S. Torchiaro, S. Tucci et al. 2013. "EAST-ADL: An architecture description language for automotive software-intensive systems," in Embedded Computing Systems: Applications, Optimization, and Advanced Design. IGI Global, pp. 456-470.
- [13] M. Bergler, J.-P. Tolvanen, M. Zoppelt, and R. Tavakoli Kolagari. 2021. "Social Engineering Exploits in Automotive Software Security: Modeling Human targeted Attacks with SAM," 31st European Safety and Reliability Conference, ESREL 2021, Sep. 2021, pp. 2502-2509, https://dx.doi.org/10.3850/978-981-18-2016-8_720-cd
- [14] MetaCase. 2023. MetaEdit+ 5.5 User's Guides, <https://metacase.com/support/55/manuals/> (accessed May 2024)
- [15] J.-P. Tolvanen and S. Kelly. 2023. "Effort used to create domain-specific modeling languages," 21st ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, Oct. 2023, <https://doi.org/10.1145/3239372.3239410>.
- [16] C. Jakobs, M. Werner, K. Schmidt and G. Hansch. 2022. "Heuristic Risk Treatment for ISO/SAE 21434 Development Projects," 2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS), Sofia, Bulgaria, pp. 653-662, <https://doi.org/10.15439/2022F136>.
- [17] M. Zoppelt and R. Tavakoli Kolagari. 2019. "SAM: A security abstraction model for automotive software systems," in Security and Safety Interplay of Intelligent Software Systems, B. Hamid, B. Gallina, A. Shabtai, Y. Elovici, and J. Garcia-Alfaro, Eds. Cham: Springer International Publishing, pp. 59-74, <https://doi.org/10.1007/978-3-030-16874-2>
- [18] M. Bergler and R. Tavakoli-Kolagari. 2023. "Automotive Software Security Engineering based on the ISO 21434", in Proceedings of the 2023 5th World Symposium on Software Engineering. Association for Computing Machinery, New York, NY, USA, 17-26, <https://doi.org/10.1145/3631991.3631994>
- [19] J.-P. Tolvanen and S. Kelly. 2023. "Evaluating Tool Support for Co-Evolution of Modeling Languages, Tools and Models", 2023 ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C), <https://doi.org/10.1109/models-c59198.2023.00144>
- [20] Macher, C. Schmittner, O. Veledar, and E. Brenner. 2020. "ISO/SAE DIS 21434 Automotive Cybersecurity Standard - In a Nutshell," Computer Safety, Reliability, and Security. SAFECOMP 2020 Workshops, pp. 123-135, <https://doi.org/10.1007/978-3-030-55583-2>

Goliath, a Programming Exercises Generator Supported by AI

Tiago Carvalho Freitas
ALGORITMI Research Centre / LASI, DI
University of Minho, Braga, Portugal
Email: tiago10cf@hotmail.com

Alvaro Costa Neto
0000-0003-1861-3545
ALGORITMI Research Centre / LASI, DI
University of Minho, Braga, Portugal

Instituto Federal de Educação,
Ciência e Tecnologia de São Paulo
Barretos, Brazil
Email: alvaro@ifsp.edu.br

Maria João Varanda Pereira
0000-0001-6323-0071
Research Centre in Digitalization and Intelligent Robotics (CeDRI)
Laboratório para a Sustentabilidade e Tecnologia
em Regiões de Montanha (SusTEC)
Instituto Politécnico de Bragança
Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Email: mjoao@ipb.pt

Pedro Rangel Henriques
0000-0002-3208-0207
ALGORITMI Research Centre / LASI, DI
University of Minho, Braga, Portugal
Email: prh@di.uminho.pt

Abstract—The teaching-learning process is complex in nature, requiring many tasks and skills to achieve success in the construction of knowledge. As per any particular kind of cognitive development, teaching and learning Computer Programming is no different in this regard: tasks must be executed, sometimes repeatedly, and skills must be developed. Despite different approaches and methodologies, exercising what has been studied is proven to be effective in pretty much any teaching-learning process. Many tools have been developed throughout time to aid in the execution of this important task, sometimes approaching the problem from the students' perspective, sometimes from the teachers'. This paper presents Goliath, a semi-automatic generator of Computer Programming exercises, whose functionality is based on Artificial Intelligence (AI) models, a Domain-Specific Language (DSL), and an online application that binds them together. Goliath's goals are directed towards teachers (and indirectly, students) by aiming to lower the burden of repeatedly constructing exercises. This is achieved through the use of templates that allow for automatic variations of an exercise to be created instantly, while relying on a common foundation. Goliath is meant to be a facilitator, raising availability of exercise lists, while avoiding repetition and the common mistakes that accompany their construction.

Index Terms—Computer Programming, Programming Education, Artificial Intelligence, Domain-Specific Languages, Programming Exercises

I. INTRODUCTION

TEACHING and learning Computer Programming is an advanced, arduous and complex process, both for teachers

This work has been supported by FCT – Fundação para a Ciência e Tecnologia within the R&D Units Project Scope: UIDB/00319/2020.

The work of Maria João was supported by national funds through FCT/MCTES (PIDDAC): CeDRI, UIDB/05757/2020 (DOI: 10.54499/UIDB/05757/2020) and UIDP/05757/2020 (DOI: 10.54499/UIDP/05757/2020); SusTEC, LA/P/0007/2020 (DOI: 10.54499/LA/P/0007/2020).

and students. Several challenges present themselves continuously, ranging from technical to personal [1], [2]. Among the many strategies to aid in this endeavour that have been researched for decades [3], [4], [5], exercising is paramount. Not only it is an opportunity for the student to grow and validate his or her lessons, but also serves as a guiding metric for the teacher, who can evaluate how to better pace the course content, and which points have been better cemented in the students' minds. In this regard, repetition is essential for a continuous and uneventful evolution in students' growth.

In order to achieve consistency in repetition, students should be able to access, answer and obtain response feedback from exercises as fast, and as frequent, as possible. In this regard, dedicated applications to aid both teachers and students in achieving the ultimate goal of teaching and learning how to program are invaluable, and have been constructed for many years. Goliath is an application that aims to leverage AI and a dedicated DSL to support teachers in creating exercises that offer students variation, availability and consistency in their training time.

This paper is further divided into six more sections. Section II contextualizes resources for practicing computer programming. Section III presents how exercises are structurally constructed, based on foundational research. Section IV compares AI-supported methods for automatic and semi-automatic generation of Computer Programming exercises. Section V details structural and functional aspects of Goliath. Section VI presents the tests, results and feedback obtained for Goliath. Finally, Section VII concludes the paper with final regards on Goliath's goals and achievements, and suggests future derivations and improvements within the scope of this research.

II. RESOURCES FOR PRACTICING PROGRAMMING

Practicing is paramount to learn and construct knowledge. It is so intrinsic to the educational process that several approaches have been developed and applied to teach and learn.

When it comes to computer programming, practice support assumes a wide range of implementations. Lists of exercises (written and printed by teachers or tutors) and problems in textbooks are classical offerings in educational contexts. More modern tools, such as online guides to programming learning [6], online courses, program animation applications [7], and automatic evaluation tools [8].

Nonetheless, creating exercises is still a challenge. Many factors demand consideration in order to create clear, useful ones: the approach to the topic under practice, the difficulty it will present to a diverse student population, the correct and unambiguous wording of the problem statement to avoid misdirection and general confusion, *etc.* Once all of these challenges have been surpassed and a collection of quality exercises is reached, creating new ones, or even variations of those that already exist, is not trivial. Besides being a time-consuming task, it becomes ever more prone to errors as repetition allows for lost of focus to creep in, resulting in typos, missed information, and incoherence. Systems to automatize these tasks have been implemented, such as SIETTE [9], which allows the creation and management of exercises repositories, and R/exams [10], a package for the R language that provides mechanisms to create both HTML and \LaTeX versions of parameterized exercise lists. Although successful in their own ways, the initial generation of the exercise components (more on that on Section III) is done directly by the user.

Goliath aims to aid in the assignment of programming exercises and their answering, but more importantly, it supports their actual construction with the use of two Large Language Models (LLMs) and a DSL-based template system, avoiding those typical mistakes and improving the possibility of variations in already proven exercises.

III. STRUCTURE OF PROGRAMMING EXERCISES

Understanding how programming exercises are designed and structured was paramount to automate parts of their construction—which is, ultimately, Goliath’s central objective. Generally speaking, these exercises were classified into different *types* that are commonly found in tests, lists, websites, *etc.* Furthermore, any of these types can be segmented into three main *components* (the statement, the code, and the answer area), each with its own responsibility in communicating the exercise’s intent to the student.

A. Exercise Types

From many definitions of exercise types available in the literature [11], [12], [13], Goliath relied on those published in [14], given their resemblance to how it handles their construction and which types it is capable of generating.

In total, seven types were collected:

- 1) **Code from scratch:** students must write down the complete solution to a problem from scratch (as the name

implies). No support code (or template) is provided, only a dedicated empty space for the answer;

- 2) **Code completion:** In order to solve this type of exercise, students must fill blanks that have been strategically positioned in a provided excerpt of code;
- 3) **Code improvement:** after being provided with a complete snippet of code that solves a given problem, students are asked to improve it. The modifications may require improvements to the performance, reduction of lines of code, use of specific constructs, *etc.*
- 4) **Bug finding:** as the name suggests, students must identify bugs (and their characteristics) in a excerpt of code, without actually correcting them;
- 5) **Debugging and fixing:** consists in a mix of the last two types, in which students are to write a correct version of a source code that contains bugs;
- 6) **Code interpretation:** students are required to interpret a given snippet of code and report on its behaviour, goals, evolution of a variable’s value throughout execution, *etc.*
- 7) **Output or state prediction:** students are asked to find out either the output of a source code’s execution, or the value of a variable throughout its lifetime.

These types are commonplace in programming courses and are implemented either physically (via printed paper, for example), or digitally. Several adaptations are also possible, including the option to transform the answer format from open to multiple choices.

Given certain implementation requirements (which will be presented and discussed later in the article), of all seven types, three were chosen and adapted to be generated by Goliath: *code from scratch*, *code completion*, and *output or state prediction*.

B. Exercise Components

There are usually three main components to consider in a typical exercise (of any of the seven types): the *problem statement*, the accompanying *code*, and the *answer area* (see Fig. 1).

The *problem statement* contains text that is presented to the student explaining the context and parameters of the problem, the type of answer that is expected, and other pertinent details about the exercise. An excerpt of *code* usually follows, containing entire programs, snippets with blanks to be filled, or concurrent versions to be compared, analysed or fixed. It supports the *problem statement* to establish a basis for solving the exercise. Finally, the *answer area* contains wither a blank space or the distribution of possible options for the student’s answer.

The definition of these three components were required to design the semi-automatic generation of exercises. Goliath based its generative mechanisms on a divide-and-conquer strategy: each component, albeit semantically connected, could be created independently, as long as the reasoning behind the problem was consistently maintained. The next step in implementing Goliath’s semi-automatic generation mechanism

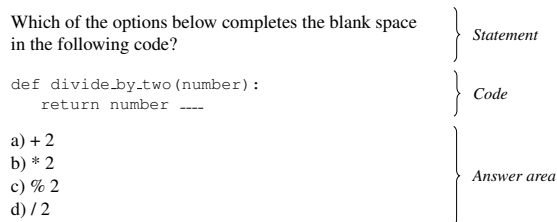


Fig. 1. Typical components of a programming exercise.

was to evaluate and choose AI models that could generate one or more of these components.

IV. AI-SUPPORTED EXERCISE GENERATION

Commonly applied to create chat-bots, translators and similar tools, Natural Language Generation (NLG) (a sub-field of Natural Language Processing (NLP)) contains theories and techniques that allow for the production of coherent and useful text in multiple languages [15], [16], including complex articles and stories [17]. The process involves taking input data, such as keywords, a set of facts or a starting piece of text, and transforming (or complementing) it into meaningful output text.

NLG is based on tasks [16] and architectures [15] that are specifically designed to produce useful text. Their features and functionalities are implemented into *models* that are trained to perform text generation, realising the foundation on which AI-supported systems are built.

A. Models

Models are pre-trained pieces of software that can recognise patterns or generate expected outputs given a set of inputs. Specifically in the case of text generation, AI models are trained to generate text based on an input of some form (keywords, sentences, questions *etc.*) [18].

The implementation techniques for AI models are varied, and have been evolving in a fast pace. Artificial Neural Networks (ANNs), a subset of Machine Learning (ML) technologies, are currently the *de facto* choice for implementing models. Their structures have been inspired by the human brain and operate in a similar fashion, offering results that are human-like [19]. Architectures of ANN that have been commonly and effectively used for text generation include:

- **Recurrent Neural Network (RNN):** suited for processing sequential data, such as text, incorporates feedback connections that take into account previous time steps and observations [20]. Implementations include Long Short-Term Memory (LSTM) networks [21], [22], and Seq2seq models [23], [24];
- **Transformers:** a ANN architecture that is capable of modelling long-range dependencies between words in a piece of text. Includes self-attention mechanisms to weight the relevance of words (or tokens) on the input. Transformers are well suited for applications that involve comprehension of context or semantics [25], [26]. Implementations include Generative Pre-training Transformer

(GPT) [27], [28], [29], [30], Bidirectional Encoder Representations from Transformers (BERT) [31], [32], and others;

- **Generative Adversarial Network (GAN):** consists of two networks linked together (a generator and a discriminator) [33], but trained individually. The primer generates text that is then verified by the later. TextGAN [34] implements and improves a GAN to generate coherent text samples;
- **Variational Autoencoder (VAE):** can be used to several goals, including text generation. It consists of two networks (an encoder and a decoder) and a latent space [35]. These three components work by trying to minimise differences between the input and a reconstructed counterpart. Bowman *et al* [36] implemented a VAE-based model to generate natural language sentences based on comparisons in two different languages.

These architectures are implemented through several techniques, languages and libraries. Examples include: OpenAI's libraries [37], TensorFlow [38], Keras [39], PyTorch [40], Hugging Face [41] and TextGenrnn [42].

B. AI Models to Generate Exercises

In order to determine which AI models would be used for the generation of the exercise components, four candidates were tested and evaluated: GPT-3.5¹, GPT-2, KeyToText [43] and CodeT5 [44]. The overall testing methodology, results and conclusions have been published in greater detail in a previous study [45]. Table I presents a resumed version of the models' features, as evaluated for the purposes of generating full programming exercises (statement, code and answer).

OpenAI's GPT-3.5 was the first model to be tested (see Listing 1 for an example). Despite returning accurate results, this version had a few shortcomings that made it less valuable for use in Goliath:

- It was free to use only within a gracing period;
- The input prompt had to contain a complete description of the exercise and how it should be generate, contradicting one of Goliath's main goals: to automate the generation of the exercise's statement;
- The output contained both the problem statement and the answer, which would require some interpretation and post-processing in order to establish an internal representation of the exercise's components;
- Even if OpenAI released it publicly (which they don't), it would be unfeasible to use it locally given its size. This means Goliath would be dependent on OpenAI's servers and consequent communication.

The second tested model was a previous version of OpenAI's offerings, GPT-2. Since it has a more manageable size, the whole model was downloaded and fine-tuned with a dataset of 300 basic Python exercises [46]. The prompt input consisted in a formatted request for the code of a specific exercise. It used Application Programming Interfaces (APIs)

¹GPT-3.5 was the newest version at the time of the study.

TABLE I
COMPARISON BETWEEN DIFFERENT AI MODELS AND THEIR FEATURES FOR PROGRAMMING EXERCISES GENERATION.

Model	Cost	Hosting	Input	Output
GPT-3.5	Paid after gracing period	OpenAI Servers	Full request and definitions	Statement, code and answer
GPT-2	Free	Local	Formatted text and statement	Statement and code
KeysToText	Free	Local	Keywords	Statement
CodeT5	Free	Local	Function statement	Code

Listing 1 Example of a programming exercise generation using GPT-3.5.

Source:

```
prompt = "Create a programming exercise in Python
about string manipulation and write the answer."
exercise = openai.Completion.create(engine="text-
davinci-003", prompt=prompt, temperature=0.7,
max_tokens=500, top_p=1, frequency_penalty=0,
presence_penalty=0)
```

Output:

Write a program to print out the last three characters of a string.

```
#Answer:
string = input("Enter a string: ")
length = len(string)
last_three_chars = string[length - 3:]
print(last_three_chars)
```

from *aitextgen* [47], a free library that interfaces with GPT-2 (see Listing 2). The output contained both the statement and the code for the exercise. It resulted in less accurate output than GPT-3.5, and some of the same shortcomings were also present (necessity to provide most of the exercise as input, and full length output that would require post-processing).

Finally, two other models were tested in tandem: KeyToText and CodeT5. KeyToText is based on Text-to-Text Transfer Transformer (T5) [48] and was fine-tuned with the Mostly Basic Python Problems (MBPP) dataset from Google Research [49], containing 1000 entry-level programming problems. The model, used with the support of the Natural Language Toolkit (NLTK) [50] platform, was able to generate accurate (albeit simple) exercise statements from a minimum of three keywords (see Listing 3 for an example). This solution was unique when compared to the GPT models, since it was able to automate the generation of the exercise statement. This statement, in turn, could be adapted and fed into the CodeT5 model to generate its accompanying source code (see Listing 4 for an example). CodeT5 also implements a T5 model, trained with the CodeSearchNet [51] dataset collection. It was able to accurately generate Python source code coherent with the exercise statement.

After approximately 100 tests were conducted, 80% rate of success was achieved in generating comprehensible and coherent exercises via KeyToText and CodeT5. Goliath implemented both models in a sequence of execution (further

Listing 2 Example of a programming exercise generation using the fine-tuned GPT-2 model.

Source:

```
prompt = """Exercise
Instructions:
Write a Python program to reverse a string.
Code:"""
model.generate(prompt=prompt, max_length=500,
temperature=0.3, top_p=0.9)
```

Output:

```
Exercise
Instructions:
Write a Python program to reverse the binary
representation of a given integer.
Code:
def test(n):
    return int(bin(n)[::-1][:-2], 2)
```

Listing 3 Example of an exercise statement generation using the fine-tuned *Keys-To-Text* model.

Source:

```
model.predict(["remove", "list", "odd numbers"])
```

Output:

Write a function to remove odd numbers from a list.

Listing 4 Example of code generation using *CodeT5*.

Source:

```
p = "Function to remove odd numbers from a list."
ids = tokenizer(p, return_tensors="pt").input_ids
code = model.generate(ids, max_length=128)
tokenizer.decode(code[0], skip_special_tokens=True)
```

Output:

```
def remove_odd_numbers(nums):
    return [n for n in nums if n % 2 == 0]
```

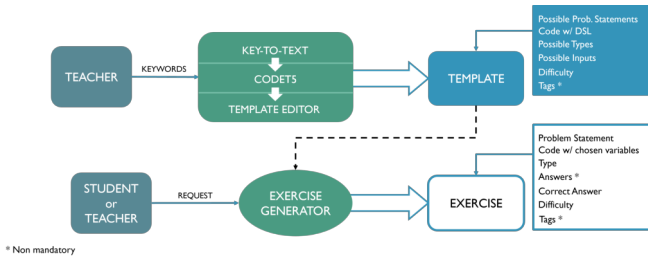


Fig. 2. Goliath's functionality and main features.

explained in Subsection V-B) to generate initial templates for Python programming exercises. This resulted in the effective use of the NLG models to aid in the construction of the whole system.

V. GOLIATH

Goliath is an online application² for semi-automatic generation of programming exercises³. It is based on two AI models (KeyToText and CodeT5), a template system, a DSL, and other supporting functionalities. It allows teachers to create a repository of templates that are used to generate programming exercises. These are then assigned to students to aid in their programming practices.

Fig. 2 shows the general design of Goliath's functionality and its main features. There two basic workflows: the creation of templates (top half of Fig. 2) and the generation of exercises (bottom half).

The template creation is based on a sequence of operations:

- 1) The teacher provides Goliath with (at least three) keywords to generate an exercise statement using KeyToText (left side of Fig. 3);
- 2) The statement, after reviewed by the teacher, is fed into CodeT5 to generate its accompanying code (right side of Fig. 3);
- 3) Both statement and code are presented to the teacher in a Template Editor;
- 4) The teacher embeds parameters into the template using commands from a DSL, specifying variations on the exercise;
- 5) The template is stored in a repository, along with a few other settings that the teacher can define (explained later in the article).

Given that the teacher may not want to use the AI-supported functionalities, both steps 1 and 2 can be skipped, as long as he or she writes the statement and the code from scratch. Further details about the template system and the DSL are presented in Subsection V-A.

²Accessible at: <https://goliath.epl.di.uminho.pt/>

³Only Python is currently supported as the programming language for the exercises. This technical limitation was implemented due to two reasons: the AI models have been fine-tuned with datasets of Python source code, and internal verification mechanisms also assume Python as the language of choice for the exercises.




Fig. 3. AI-supported statement and code generation.

Creating a template does not generate an exercise in itself. This only happens after a request is sent to the Exercise Generator by the student. The generation process is straightforward:

- 1) The student requests an exercise that was assigned by the teacher;
- 2) The corresponding template is fetch from the repository;
- 3) A version of the exercise is automatically generated by choosing randomly one of the pre-defined variations that the template specifies;
- 4) The answer alternatives are created *ad hoc*, based on the exercise version;
- 5) The exercise (statement, code and answer alternatives) is presented to the student who can answer it;
- 6) The correct answer is shown as feedback.

This division of workflows is intentional: by delaying the generation of the exercise, Goliath is able to add a layer of controlled randomisation to the process. This fact contributes to its replayability and provides a sense of discovery to the students. Further explanations on the exercise generation are given in Subsection V-B.

A. Goliath's DSL and the Template System

Goliath's template system allows educators to create *replayability*⁴, in which one template can automatically generate different versions of an exercise.

Versions, in this context, represent variations of what is requested from students. As an example, a simple exercise statement could read "Write a function that removes all odd numbers from a list of integers". In order to change the request from *remove odd numbers* to *remove even numbers*, this statement needs only to be minutely changed. Furthermore, the statement could also be modified to request the removal of all positive numbers, all prime numbers, or any equivalent variation. Thus, if given the possibility of automatically generating these versions (odd, even, positive, prime, *etc.*) from one template, students could practice multiple times, while educators needed only to construct and parameterise the template once.

The template consists in the statement, the code, and a few added settings (explained in Subection V-C). The answer area is the only component that is not directly included in the template, as it is automatically created when the exercise is generated. For demonstration purposes, consider Listing 5 as a basis for a template that was suggested by the AI models.

⁴Replayability indicates the possibility of using the system multiples with a lower chance of encountering the same state with frequency.

Listing 5 Basis for a template created from the keywords *list*, *remove*, and *numbers*.

Statement:

Function that removes all odd numbers from a list of integers.

Code:

```
def remove_odd_numbers(nums):
    return [n for n in nums if n % 2 == 0]
```

Listing 6 Grammar for the DSL that allows templates to generate several versions of an exercise. Note that the commands can not be recursively nested.

```
start          : declaration* commands
declaration   : ID "=" STRING ("," STRING)+
commands      : (alternatives | conditional)*

alternatives: STRING ("," STRING)* fail?
conditional : "case" vars ":" actions else? fail?
vars        : ID ("," ID)*
actions     : action (";" action)*
action      : STRING ("," STRING)* "=>" STRING

else        : ";" "else" ":" STRING
fail       : ";" "fail" ":" STRING ("," STRING)*
```

If left unchanged, this template would only generate one version of the exercise (“*Function that removes all odd numbers from a list of integers*”). In order to make it more flexible, it is necessary to parameterise it. A DSL was designed specifically for this purpose. Its grammar is presented in Listing 6, which was used with Lark [52] to develop an interpreter. This DSL is intended to be used in the same fashion as a markup language, whose commands must be interspersed into the text.

The commands of parameterisation are placeholders delimited by double curly brackets (`{{ and }}`). They are substituted by a value (generally speaking, a piece of text) when the exercise is generated. There are three types of commands: *key declarations*, *conditional placements*, and *simple alternatives*.

1) *Key declaration*: Multiple keys can be declared in the exercise statement, representing variations on what is requested from the student. Each key defines a list of possible values. For example, in order to parameterise the template of Listing 5 and allow the generation of both *odd* and *even* versions of its exercise, the statement should be rewritten as:

```
Function that removes all {{ x = 'odd', 'even' }}
numbers from a list of integers.
```

The key `x` creates the variation for *odd* and *even*. When an exercise is generated from this template, its statement will read either “*...function that removes odd numbers...*” or “*...function that removes even numbers...*”. Listing 7 shows the template with the parameterised statement.

Although unnecessary for this example, there could be multiple key declarations. The following statement would be

Listing 7 Template with a key declaration in the statement.

Statement:

Function that removes all {{ x = 'odd', 'even' }} numbers from a list of integers.

Code:

```
def remove_odd_numbers(nums):
    return [n for n in nums if n % 2 == 0]
```

able to generate four different versions of the exercise (remove *odd* from a *list*, remove *odd* from an *array*, remove *even* from a *list*, and remove *even* from an *array*):

```
Function that removes all {{ x = 'odd', 'even' }}
numbers from {{ y = 'a list', 'an array' }} of
integers.
```

Finally, key declarations must only occur in the statement, as they establish variations on what the exercise will ask from the student. They would not make sense in the code, given that its logic is supposed to be derived from the statement, not the other way around.

2) *Conditional placements*: Assuming that an exercise must be entirely coherent, if there are possible variations on what is asked from student (via *key declarations* in the statement), a fixed code would probably be wrong. *Conditional placements* allow the code to adapt to the variations of the keys when the exercise is generated.

Conditional placements always start with the keyword `case`, followed by the list of keys that should be considered. For each combination of values, a result is specified using the *hash-rocket* notation (`=>`). It is similar to a `switch-case` in a conventional programming language.

Going back to the example in Listing 7, if the key `x` assumes the value `odd` in the exercise, the second line of code must read:

```
return [n for n in nums if n % 2 == 0]
```

This guarantees that only even numbers are left in the list, effectively removing odd numbers. On the other hand, if `even` is chosen for `x`, the line must change to:

```
return [n for n in nums if n % 2 == 1]
```

The template can be parameterised for this variation through a conditional placement on `x`:

```
return [n for n in nums if n % 2 == {{ case x: 'odd'
=> '0'; 'even' => '1' }}]
```

This conditional placement results⁵ in 0 if `x` is `odd`, and 1 if `x` is `even`. Since `x` has only two possible values, the command could also be written using an `else` clause:

```
return [n for n in nums if n % 2 == {{ case x: 'odd'
=> '0'; else: '1' }}]
```

⁵As previously explained, all commands from the DSL are placeholders. The result of a command refers to the value that will substituted it in the exercise.

This clause acts as a complementary branch in the conditional placement. It is similar to the default branch in a switch-case.

Finally, the `fail` clause can be used to specify invalid options for the placement:

```
return [n for n in nums if n % 2 == {{ case x: 'odd'
=> '0'; else: '1'; fail: '2', '3' }}]
```

The definition of invalid options using the `fail` clause is important as Goliath is unable to automatically come up with wrong answers (further explanations in Subsection V-B).

Finally, a conditional placement can take multiple keys into consideration:

```
{{ case x, y: 'odd', 'a list' => '0';
'even', 'an array' => '1';
else: '-1';
fail: '2', '3' }}
```

3) *Simple alternatives*: Unlike conditional placements, *simple alternatives* do not take keys into consideration. They simply create variations, with the added possibility of specifying invalid options:

```
return [n for n in nums if n {{ '% 2'; fail: '/ 2',
'* 2', '+ 2' }} == 0]
```

The correct option for that portion of code is `% 2`, while the others (defined by the `fail` clause) are invalid and will generate incorrect alternatives for the answer in the exercise.

Simple alternatives can also be used to create equivalent variations in the code:

```
h = {{ 'n / 2.0', 'n * 0.5' }}
```

Both variations of the code above are equivalent in their objective (assigning the half of `n` to `h`) and would result in two different, albeit valid, versions.

4) *Final remarks on the DSL*: The three types of commands (key declarations, conditional placements and simple alternatives) allow for the definition of both valid and invalid variations. Both are equally important for the generation of the exercise, given that Goliath is not able to self-determine which changes in the code would result in correct answers and which would not. In order to that, Goliath would be required to not only perceive what is asked by the statement—meaning, the actual problem that the exercise entails—but also to check if the code’s logic is coherent with it. The Halting Problem guarantees this impossibility.

Typical error checking is done during the interpretation of the DSL commands, including the use of undeclared keys, syntactic mistakes and invalid commands. After all commands have been successfully interpreted, an internal JSON-like representation of the keys and their variations is created, and stored with the template (see Listing 8).

Listing 9 shows the complete version of the exercise template, including all three types of commands.

B. Exercise Generation

As previously mentioned, a template acts as a blueprint that generates different versions of an exercise. The generative

Listing 8 Internal representation of a template’s keys and variations.

```
"keys": {
  "x": ["odd", "even"]
},
"variables": {
  "alt1": {
    "correct": [
      { "value": "% 2" }
    ],
    "wrong": ["/ 2", "* 2", "+ 2"]
  },
  "alt2": {
    "correct": [
      {
        "value": "0",
        "conditions": [
          { "key": "x", "index": 0 }
        ]
      },
      {
        "value": "1",
        "conditions": [
          { "key": "x", "index": 1 }
        ]
      }
    ],
    "wrong": ["2", "3"]
  }
}
```

process begins with a request from a student, when he or she accesses an assignment created by the teacher⁶. In this context, the real exercise can be considered an instantiation of the assignment’s template.

A few important considerations must be made in order to understand the exercise generation:

- There are three types of exercises available in Goliath: *code selection* (adapted from *code from scratch*), *input/output*, and *code completion* (also adapted)⁷. This limitation was imposed by the mechanisms implemented in the exercise generation. Additional types would require more commands in the DSL and more settings in the template, which was unfeasible for this version of Goliath;
- Answers are presented in multiple choice format for the *code selection* and *code completion* types⁸. Implementing an open answer format for these types would either negate the immediate feedback to students, as they waited for the teacher to correct the answers, or require the implementation of an extremely accurate NLP model. *input/output*, on the other hand, is presented in open answer format, since it can be trivially verified by executing the code of the exercise itself;
- Both the statement and the code in the template need to follow specific discourses. Statements must start with

⁶Teachers may create an assignment at any time, as long as the template has already been created and stored in the repository.

⁷See Subsection III-A for the description of each type.

⁸This is the reason Goliath implements adapted versions of the original exercise types.

Listing 9 Template containing all three constructs.**Statement:**

Function that removes all `{{ x = 'odd', 'even' }}` numbers from a list of integers.

Code:

```
def remove_{{ x }}_numbers(nums):
    return [n for n in nums if n {{ '% 2'; fail: '/ 2', '* 2', '+ 2' }} == {{ case x: 'odd' => '0';
                                                                    else: '1';
                                                                    fail: '2', '3' }}]
```

“Function to...”, “Function that...”, etc. while the code should always contain a single function definition. This requirement exists because Goliath needs to adapt them to the exercise type when the generation occurs, which requires complementing the statement text and running the code’s function.

Goliath follows a simple routine to generate an exercise:

- 1) The type of the exercise is randomly picked within the range of possible options;
- 2) The value for each key in the statement is randomly chosen;
- 3) The code is adjusted to the values of the keys;
- 4) The answer alternatives are calculated;
- 5) The exercise is constructed from the three components (statement, code and answer alternatives).

The calculation of the answer alternatives demands further explanation. The template carries more than just the statement and the code. It also contains other settings that are used to determine the possible exercise types, the answer alternatives, its correctness, difficulty and category (more on this in Subsection V-C). One of these settings is a list of valid inputs for the function defined in the code. This list is used to generate exercises of the *input/output* type, which read “*What is the output of the following function when the input is ...?*”. It is also used to check the student’s answer by comparing it to the output of the function when fed with the valid input.

The incorrect alternatives for the other two types of exercises are constructed using the values of the `fail` clauses in the code. In the end, the presence or absence of valid inputs and `fail` clauses determine which exercise types can be generated. Listing 10 shows one example for each type of exercise based on the template of Listing 9. In these examples, even was chosen for the key `x` of the statement.

C. Goliath’s Interface

Goliath was implemented as an online application using a mix of technologies and languages (a Programming Cocktail) containing Python, Flask, Lark and MongoDB. Its interface follows the general workflows presented in Fig. 2, with a few added pages to manage users, control access to the several parts of the application, define assignments, and manage templates and exercises.

Listing 10 Three exercise types generated from the same template.**Code selection:**

Which of these options is a function that removes all even numbers from a list of integers.

- a) `def remove_even_numbers(nums):`
 `return [n for n in nums if n % 2 == 0]`
- b) `def remove_even_numbers(nums):`
 `return [n for n in nums if n % 2 == 1]`
- c) `def remove_even_numbers(nums):`
 `return [n for n in nums if n * 2 == 0]`
- d) `def remove_even_numbers(nums):x`
 `return [n for n in nums if n % 2 == 3]`

Input/output:

What is the output of the following function when the input is [1, 2, 3, 4, 5].

```
def remove_even_numbers(nums):
    return [n for n in nums if n % 2 == 1]
```

Answer: _____

Code completion:

Which of these options complete the following function that removes all even numbers from a list of integers.

```
def remove_even_numbers(nums):
    return [n for n in nums if n % 2 == ____]
```

- a) 1
- b) 0
- c) 2
- d) 3

The main pages of the application are the *AI suggestion page*, the *template edit form*, the *exercises management page* and the *exercise page*.

The AI suggestion page (Fig. 4) is the starting point to the creation of a new template in which the teacher inputs keywords for the KeyToText model to process and suggest a statement. After review (the lower input field), the teacher can send this statement for the CodeT5 to generate the associated code and fill the next page, the template edit form.

The template edit form (Fig. 5) allows the teacher to edit

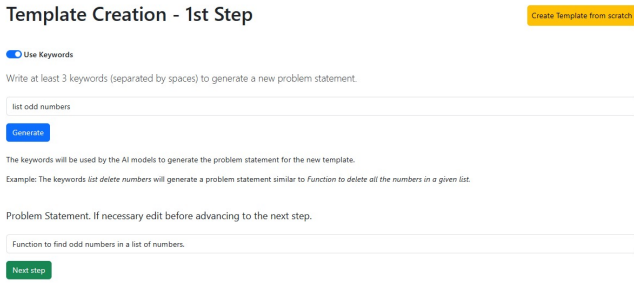


Fig. 4. Keyword input for the AI model.

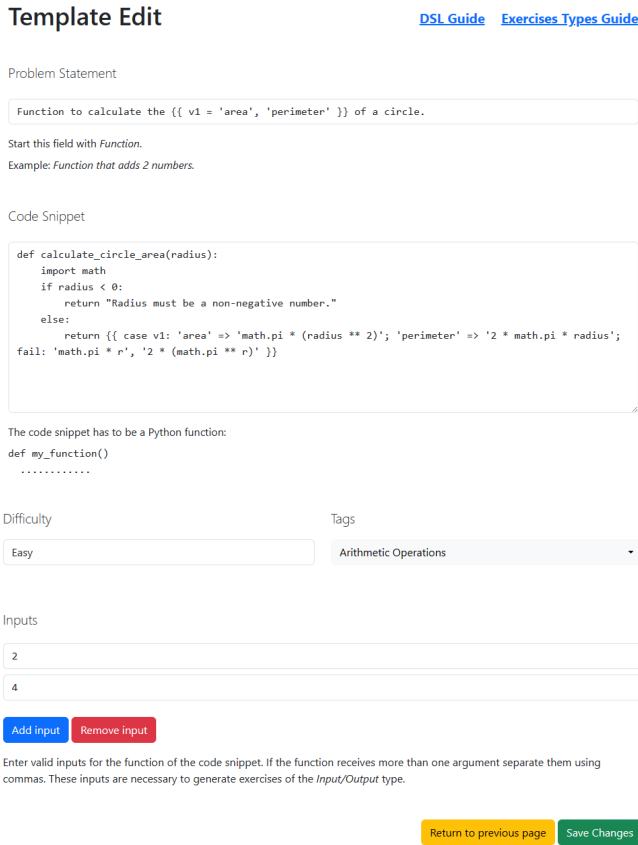


Fig. 5. Template edit form.

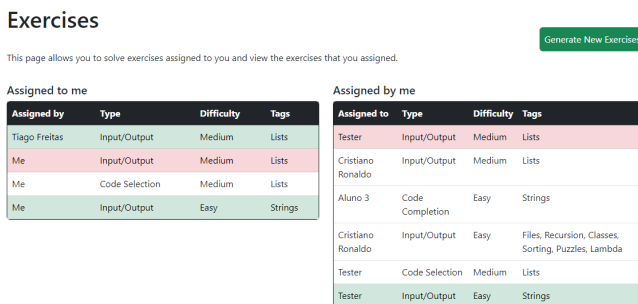


Fig. 6. Exercises management page.

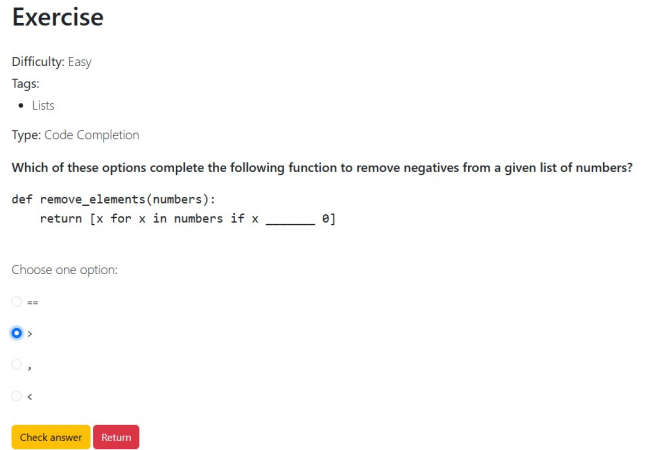


Fig. 7. Exercise page.

the statement, the code and the other settings for the template. It comes after the AI models have made their suggestions (or if they were skipped entirely). The settings for the template are:

- **Difficulty:** a difficulty level from easy to hard designed to guide students in approaching the easier exercises first. Defined entirely by the teacher;
- **Tags:** a form of categorization for the template. Specifies the general programming concepts that the exercises of this template will tackle;
- **Inputs:** a list of valid inputs for the function defined in the code. As explained at the end of Subsection V-B, these inputs will be used to generate and verify the answer to exercises of the *input/output* type.

The exercises management page (Fig. 6) has two points-of-view: the teacher has access to the right table and the *Generate New Exercise* button, while the student only sees and interacts with the left table.

The right table lists exercises that have been assigned to students and it is only visible to teachers. Each entry in the table represents one assignment and its background color indicates if the student was correct (green), incorrect (red) or still have not answered (white). The button labeled *Generate New Exercises* is used to create the assignments.

The left table is only visible for users with the student role, as they represent exercises that have been assigned to them by the teachers. In fact, the exercise requests explained in Subsection V-B effectively occurs when a student first opens entries from this table.

Finally, the exercise page (Fig. 7) allows the student to answer an exercise and obtain immediate feedback of his or her response. The image shows an exercise of type *code completion* begin answered.

VI. TESTS AND FEEDBACK

In order to evaluate Goliath’s functionality, a survey was conducted with teachers from Computer Programming back-

ground. The survey consisted in ten questions: two for establishing the respondent's background, seven to evaluate Goliath's features and usability—in a scale from 1 (terrible) to 5 (excellent)—and a final open question for additional remarks and feedback.

The questions, in order of appearance, were:

- 1) Programming experience (in years).
- 2) Experience in programming teaching (in years).
- 3) Regarding general usability, how do you rate the ease of navigation and interaction with the application?
- 4) How do you rate the ease of generating text (instructions and code) from the AI models?
- 5) How do you rate the quality of the text generation results (instructions and code) of the AI models?
- 6) How do you rate the way the templates are structured?
- 7) How do you rate the influence of the DSL in generating different exercises from the same template?
- 8) How do you rate the ease of using the DSL?
- 9) How do you rate the quality of the generated exercises?
- 10) In this section, you are asked to comment on any aspects not covered by the questions and to report errors/bugs that have appeared while using the application.

The survey was answered by 10 people and the results produced the charts in Fig. 8.

The first two questions revealed that the survey was answered by people with different levels of experience, including those that have never taught programming. This result indicates that the evaluation for the next seven questions are not skewed towards teachers, but a general overview of the application.

The results were mostly positive, especially in usability and ease of use of the main functionalities. The question with the lowest rating is the fifth (*Quality of the text generated by the AI models*), which indicates that the AI models have not been completely effective. Although, some respondents also indicated that the lower grades were due to instability of the models. The processes that run the models hung up the application a few times when the server was under heavier load from other sources. Nevertheless, the evaluations show that, despite a few inefficiencies, their preparation and the detailed processing of the input and output texts, resulted in average to good results.

The feedback obtained from the last question mainly centered around suggestions of features that would complement the application, such as the possibility to create entire tests inside the application, and even a layout suitable for printing. Small bugs were also reported, which were addressed, making the application as consistent as possible.

VII. CONCLUSION

This paper presented Goliath, an online application aimed at supporting teachers and students in the practice of Computer Programming. It is based on two AI models to kickstart the construction of programming exercise templates. The templates are parameterised to generate different versions of an exercise. This is done through commands of a DSL that was

developed specifically for Goliath. Exercises are generated *ad hoc* when a students request them via assignments created by their teachers. Also, the exercises may be requested by the teachers and used indirectly in their classes or to compose offline lists.

Through a testing period and a survey, results showed that Goliath is already in a working state, capable of supporting teachers in their educational endeavours. A few facets of the User Experience can be improved in order to allow for more efficiency and efficacy in the resulting exercises. Overall, Goliath fulfilled its foundational goal of using AI models in a supportive way, while also providing teachers with high level flexibility and control in the entire process.

Suggestions for future works include both new features and improvements to the existing ones. Among the new features, new modules to apply Goliath to tests and other practice-oriented situations would be beneficial for teachers, supervisors and tutors. Also, some mechanisms to provide greater independence to students would be of great value, such as the automatic generation of complete lists of exercises based on their history and background. This features would also free teachers from the assignment task, which could stimulate Goliath's adoption in the educational setting. On the improvement side, the quality of both the statement and the code generated by the AI models should be improved, in order to obtain more variety and complexity in their suggestions. Finally, reliability and efficiency in the communication between Goliath's internal parts and the models could also be improved.

REFERENCES

- [1] A. Gomes and A. J. Mendes, "Learning to program: Difficulties and solutions," Proceedings of the 2007 International Conference on Engineering and Education (ICEE). International Network on Engineering Education and Research, 2007, pp. 283–287. [Online]. Available: <http://icee2007.dei.uc.pt/proceedings/papers/411.pdf>
- [2] J. Figueiredo and F. J. García-Peñalvo, "Building skills in introductory programming," F. J. García-Peñalvo, Ed., Proceedings of the Sixth International Conference on Technological Ecosystems for Enhancing Multiculturality. New York: ACM, 10 2018. doi: 10.1145/3284179. ISBN 9781450365185 p. 46–50. [Online]. Available: <https://dl.acm.org/doi/10.1145/3284179.3284190>
- [3] M. J. V. Pereira and P. R. Henriques, "Visualization/animation of programs in alma: Obtaining different results," in *Proceedings of the IEEE Symposium on Human Centric Computing Languages and Environments*, 2003. doi: 10.1109/HCC.2003.1260242 pp. 260–262. [Online]. Available: <https://ieeexplore.ieee.org/document/1260242>
- [4] R. R. Fenichel, J. Weizenbaum, and J. C. Yochelson, "A program to teach programming," *Communications of the ACM*, vol. 13, pp. 141–146, 03 1970. doi: 10.1145/362052.362053. [Online]. Available: <https://dl.acm.org/doi/10.1145/362052.362053>
- [5] S. A. Robertson and M. P. Lee, "The application of second natural language acquisition pedagogy to the teaching of programming languages: a research agenda," *ACM SIGCSE Bulletin*, vol. 27, no. 4, p. 9–12, 12 1995. doi: 10.1145/216511. [Online]. Available: <https://dl.acm.org/doi/10.1145/216511.216517>
- [6] M. V. P. Almeida, L. M. Alves, M. J. V. Pereira, and G. A. R. Barbosa, "Easycoding: Methodology to support programming learning," R. Queirós, F. Portela, M. Pinto, and A. Simões, Eds., vol. 81, Open Access Series in Informatics (OASISs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 06 2020. doi: 10.4230/OASISs.ICPEC.2020.1. ISBN 978-3-95977-153-5. ISSN 2190-6807 pp. 1–8. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2020/12288>
- [7] Python Tutor, "Python tutor." [Online]. Available: <https://pythontutor.com>

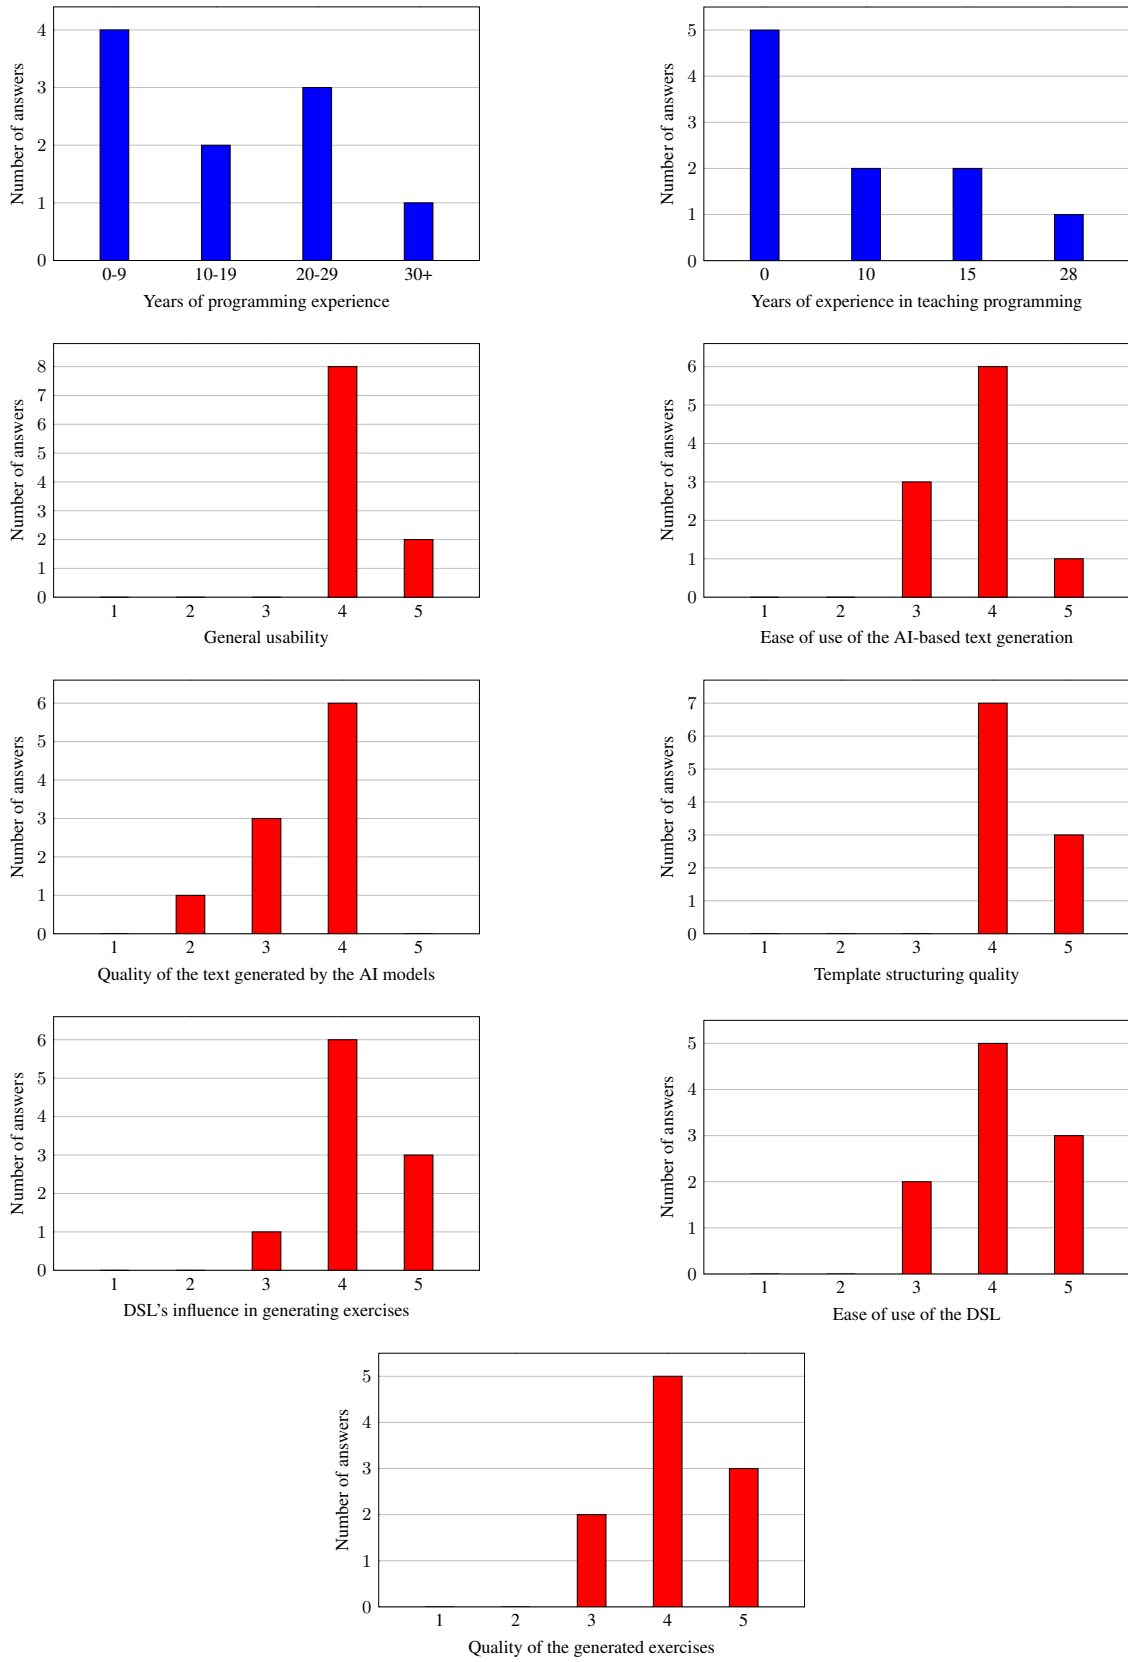


Fig. 8. Results from the survey.

- [8] beecrowd, "beecrowd," 2024. [Online]. Available: <https://beecrowd.com>
- [9] A. Rios, J. L. Pérez de la Cruz, and R. Conejo, "Siette: Intelligent evaluation system using tests for teleeducation," in *WWW-Based Tutoring Workshop at 4th International Conference on Intelligent Tutoring Systems*, 1998. [Online]. Available: <https://www.siette.org>
- [10] A. Zeileis, "R/exams." [Online]. Available: <https://www.r-exams.org>
- [11] M. Bower, "A taxonomy of task types in computing," *SIGCSE Bull.*, vol. 40, no. 3, p. 281–285, jun 2008. doi: 10.1145/1597849.1384346. [Online]. Available: <https://doi.org/10.1145/1597849.1384346>
- [12] A. Ruf, M. Berges, and P. Hubwieser, "Classification of programming tasks according to required skills and knowledge representation," vol. 9378, 09 2015. doi: 10.1007/978-3-319-25396-1_6. ISBN 978-3-319-25395-4
- [13] N. Ragonis, "Type of questions - the case of computer science," *Olympiads in Informatics*, vol. 6, pp. 115–132, 01 2012.
- [14] A. Simões and R. Queirós, "On the Nature of Programming Exercises," in *First International Computer Programming Education Conference (ICPEC 2020)*, ser. OpenAccess Series in Informatics (OASlcs), R. Queirós, F. Portela, M. Pinto, and A. Simões, Eds., vol. 81. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020. doi: 10.4230/OASlcs.ICPEC.2020.24. ISBN 978-3-95977-153-5. ISSN 2190-6807 pp. 24:1–24:9. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2020/12311>
- [15] E. Reiter and R. Dale, *Building Natural Language Generation Systems*. Cambridge University Press, 2000.
- [16] —, "Building applied natural language generation systems," *Natural Language Engineering*, vol. 3, 03 2002.
- [17] A. Celikyilmaz, E. Clark, and J. Gao, "Evaluation of text generation: A survey," *CoRR*, vol. abs/2006.14799, 2020. [Online]. Available: <https://arxiv.org/abs/2006.14799>
- [18] IBM, "What is an AI model?" 2024. [Online]. Available: <https://www.ibm.com/topics/ai-model>
- [19] —, "What are Neural Networks?" <https://www.ibm.com/topics/neural-networks>, 2023.
- [20] —, "What are Recurrent Neural Networks?" <https://www.ibm.com/topics/recurrent-neural-networks>, 2023.
- [21] A. Graves, "Generating sequences with recurrent neural networks," *CoRR*, vol. abs/1308.0850, 2013. [Online]. Available: <http://arxiv.org/abs/1308.0850>
- [22] A. Karpathy, "The unreasonable effectiveness of recurrent neural networks," 2015. [Online]. Available: <http://karpathy.github.io/2015/05/21/mneffectiveness/>
- [23] P. Dugar, "Attention — seq2seq models," <https://towardsdatascience.com/day-1-2-attention-seq2seq-models-65df3f49e263>, 2019.
- [24] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS'14. Cambridge, MA, USA: MIT Press, 2014, p. 3104–3112.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017. ISBN 9781510860964 p. 6000–6010.
- [26] IBM, "What is Unsupervised Learning?" <https://www.ibm.com/topics/unsupervised-learning>, 2023.
- [27] OpenAI, "OpenAI," <https://www.openai.com/product>, 2023.
- [28] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," *CoRR*, vol. abs/2005.14165, 2020. [Online]. Available: <https://arxiv.org/abs/2005.14165>
- [29] OpenAI, "Gpt-4 technical report," 2023.
- [30] —, "ChatGPT," <https://openai.com/blog/chatgpt>, 2023.
- [31] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019. doi: 10.18653/v1/n19-1423 pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>
- [32] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "Bertscore: Evaluating text generation with bert," 2020.
- [33] S. Ali, D. DiPaola, and C. Breazeal, "What are gans?: Introducing generative adversarial networks to middle school students," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 17, 2021. [Online]. Available: <https://par.nsf.gov/biblio/10252915>
- [34] Y. Zhang, Z. Gan, K. Fan, Z. Chen, R. Henao, D. Shen, and L. Carin, "Adversarial feature matching for text generation," 2017. [Online]. Available: <https://arxiv.org/abs/1706.03850>
- [35] T. Cemgil, S. Ghaisas, K. Dvijotham, S. Goyal, and P. Kohli, "The autoencoding variational autoencoder," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 15 077–15 087. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/ac10ff1941c540cd87c107330996f4f6-Paper.pdf
- [36] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Józefowicz, and S. Bengio, "Generating sentences from a continuous space," *CoRR*, vol. abs/1511.06349, 2015. [Online]. Available: <http://arxiv.org/abs/1511.06349>
- [37] PyPI, "OpenAI," <https://pypi.org/project/openai/>, 2023.
- [38] TensorFlow, "Why TensorFlow," <https://www.tensorflow.org/about>, 2023.
- [39] Keras, "Keras documentation: About Keras," <https://keras.io/about/>, 2023.
- [40] J. Terra, "Pytorch Vs Tensorflow vs Keras," <https://www.simplilearn.com/keras-vs-tensorflow-vs-pytorch-article>, 2023.
- [41] HuggingFace, "Transformers," <https://huggingface.co/docs/transformers/index>, 2023.
- [42] M. Wolf, "Textgenrn," <https://github.com/minimaxir/textgenrn>, 2020.
- [43] G. Bhatia, "keytotext." [Online]. Available: <https://github.com/gagan3012/keytotext>
- [44] Y. Wang, W. Wang, S. Joty, and S. C. Hoi, "CodeT5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 8696–8708. [Online]. Available: <https://aclanthology.org/2021.emnlp-main.685>
- [45] T. C. Freitas, A. Costa Neto, M. J. a. V. Pereira, and P. R. Henriques, "NLP/AI Based Techniques for Programming Exercises Generation," in *4th International Computer Programming Education Conference (ICPEC 2023)*, ser. Open Access Series in Informatics (OASlcs), R. A. Peixoto de Queirós and M. P. Teixeira Pinto, Eds., vol. 112. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi: 10.4230/OASlcs.ICPEC.2023.9. ISBN 978-3-95977-290-7. ISSN 2190-6807 pp. 9:1–9:12. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2023/18505>
- [46] w3resource, "Python exercises, practice, solution," <https://www.w3resource.com/python-exercises/>, 2023.
- [47] M. Woolf, "aitextgen," <https://docs.aitextgen.io/>, 2021.
- [48] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *J. Mach. Learn. Res.*, vol. 21, no. 1, jan 2020.
- [49] J. Austin, A. Odena, M. Nye *et al.*, "Program synthesis with large language models," *arXiv preprint arXiv:2108.07732*, 2021.
- [50] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit*. O'Reilly Media, Inc., 2009.
- [51] H. Husain, H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "Codesearchnet challenge: Evaluating the state of semantic code search," *CoRR*, vol. abs/1909.09436, 2019. [Online]. Available: <http://arxiv.org/abs/1909.09436>
- [52] E. Shinan, "Lark." [Online]. Available: <https://lark-parser.readthedocs.io/en/stable/>

IoB-TMAF: Internet of Body-based Telemedicine Adoption Framework

Taif Ghiwaa

0000-0003-2990-8605

Department of Computer Science and Information Systems,
King Khalid University, Abha, Saudi Arabia

Department of Informatics,

University of Sussex,

Brighton, United Kingdom

Email: t.ghiwaa@sussex.ac.uk

Imran Khan, Martin White, and Natalia Beloff

0000-0002-5732-4685

0000-0001-8686-2274

0000-0002-8872-7786

Department of Informatics,

University of Sussex,

Brighton, United Kingdom

Email: {imran.khan,m.white,n.beloff}@sussex.ac.uk

Abstract—Saudi healthcare organizations are increasingly using Telemedicine (TM) services to reduce expenses and improve the effectiveness of healthcare delivered. Population aging and the growth of the costs of chronic diseases management has an urgent problem that requires the use of technical solutions that contribute to expanding and improving healthcare services and addressing these issues. Consequently, the growing investments in developing TM products and services have made user acceptance of technology crucial in ensuring effective use. The purpose of this study is to explore the factors influencing Saudi patients and healthcare providers to adopt Internet of Body (IoB) technologies to support diagnosis in TM settings. The Technology Acceptance Model (TAM) is employed in this study as the foundational theoretical framework, extending it with additional constructs to fit the context. The IoB-TMAF model identifies factors influencing the adoption intentions of patients and providers for IoB-based TM system. The influencing factors stem from users' individual contexts (social influence, self-efficacy, attitude, and perceived trust), technological contexts (perceived usefulness, perceived ease of use, task fit, reliability, perceived cost, and perceived privacy control), organizational contexts (facilitating conditions), and health contexts (perceived health risk). This study adds to the existing literature by introducing a comprehensive model to explore the motivational factors driving the effective adoption of IoB-based TM in the Kingdom of Saudi Arabia (KSA). Thus, formulating a strategy for the proper execution aligned with the viewpoints of its users.

Index Terms—telemedicine, Technology Acceptance Model, adoption of telemedicine, IoB, TAM.

I. INTRODUCTION

THE information revolution has seen significant advancements in telecommunication services within the health sector, transforming how healthcare is delivered, managed, and experienced. A notable example of this shift can be clearly seen in the case of Telemedicine (TM), which delivers remote clinical services over Information and Communication Technologies (ICT) channels, such as consultation, medical diagnostic, and monitoring. The global TM market is projected to reach around \$300 billion by 2030 [1], driven by factors like the COVID-19 pandemic, workforce shortages, and the increasing prevalence of chronic diseases [1], [2], [3].

The benefits of TM have been revealed by many scientists in their literature. It promotes equitable access to quality care regardless of location or physical limitations. According to studies [4] and [5], TM is a prominent solution for bridging geographical and social barriers to healthcare by reaching rural and underserved communities, people with disabilities, the elderly, and beyond. Moreover, research has shown that TM plays a vital role in improving healthcare outcomes [6], expanding access to specialized consultations [7], and reducing travel time and expenses for patients [8]. Therefore, TM becomes a powerful tool in ensuring healthcare accessibility, affordability, and quality, which gives it a significant role to play in any healthcare system.

The potential of TM is further revolutionized by the Internet of Things (IoT) technologies. Studies have shown IoT significantly enhances TM by increasing usability, acceptance, and adoption [2], [9], [10]. This interest has led to the emergence of the Internet of Body (IoB) in 2016, which integrates the human body into interconnected systems to collect health data [11]. The IoB includes a range of different technologies like wearable, embedded medical devices, and sensors that collect health data. This innovative approach promises to personalize and enhance healthcare services further.

Given TM's significant potential, both developed and developing countries are heavily investing in remote healthcare services. TM is integral to healthcare systems in countries like the USA, UK, Canada, and the EU, where infrastructure and policies support its implementation [12], [13]. Meanwhile, developing countries, particularly in the Middle East, are building the necessary infrastructure, training healthcare professionals, and creating policies to expand TM services [13], [8].

Furthermore, the author of [13] highlighted compelling reasons for healthcare organizations in developing countries to implement TM systems. Adopting TM is critical to address the growing challenges posed by aging populations while reducing costs and maintaining high-quality healthcare delivery. Moreover, the World Health Organization (WHO) noted that technological advancements and the widespread availability of

affordable internet and smart devices have paved the way for these countries to integrate TM into their healthcare systems [14]. This has enabled healthcare organizations to offer innovative and efficient healthcare services by expanding TM's scope to include web-based applications such as email, video conferencing, and sharing medical imagery with professionals [13]. Given the multifaceted benefits of TM at various levels, it is imperative to capitalize on its potential by addressing digital divides and employing best practices to ensure successful adoption and utilization of this technology for equitable access to high-quality healthcare.

Despite TM's promise, the degree of success in TM applications varies between developed and developing countries. It has been reported that several TM pilot projects have been terminated due to the many challenges they encountered [15], [16], [17]. Additionally, studies have indicated that TM has not been fully integrated into the healthcare system to provide routine services as intended [18], [19]. There might be some factors that may contribute to this, such as legal and regulatory barriers, data privacy concerns, technological limitations, lack of physical assessments, incompatibility with medical workflows, or a lack of perceived utility for current solutions [20]. Moreover, a growing body of research has shown that individuals' intentions and perspectives toward adopting and accepting a novel system are strongly influenced by users' behaviour and their culture [15], [9], [21], [22], [23]. Therefore, user acceptance is becoming an essential factor in determining the success of IT implementation or the introduction of new systems.

While the existing literature on TM adoption is extensive, there remains a notable gap regarding the intention to use IoB as a supporting tool for diagnosis in the TM context. Additionally, to date, the factors influencing its acceptance by healthcare providers and patients in the Kingdom of Saudi Arabia (KSA) are not fully identified and understood [9]. Therefore, the current study aims to fill the gap by investigating the factors that shape the acceptance of IoB-based TM among Saudi healthcare providers and patients. To achieve this, the Technology Acceptance Model (TAM) is adapted and extended to fit the context of IoB-based TM. Understanding the factors that drive the acceptance of IoB-based TM is crucial for its successful implementation and integration into healthcare systems. The findings will provide insights into improving healthcare delivery, particularly in regions like the KSA where TM can address significant healthcare challenges.

The following sections will discuss the theoretical foundation of the study, develop the IoB-TMAF theoretical framework, and detail its constructs. The final section will conclude the discussion, outline future research directions, and propose a methodology for examining the IoB-TMAF model.

II. THEORETICAL FOUNDATION

Examination of the literature uncovers a range of theoretical models that provide insight on understanding users' intentions and motivations to adopt ICT [21]. Some examples of these models are Diffusion of Innovations Theory (DOI) [22],

Theory of Reasoned Action (TRA) [23], Theory of Planned Behaviour (TPB) [24], TAM and its extensions [25], [26], Unified Theory of Acceptance and Use of Technology (UTAUT) [27], and Health Belief Model (HBM) [28]. These models present various factors that impact the behaviour of end users in adopting IT. The significance of addressing the adoption of IoB technology in TM services, and the development of a new framework for investigating influential factors, becomes evident through a systematic review of the current literature [9]. The results of this review have indicated that two models, TAM and UTAUT, are the most common models employed to understand the factors influencing the adoption of TM among healthcare providers and patients across diverse countries and TM settings [9]. Furthermore, most studies included in the review introduce additional contextual factors and integrate them with the base models, such as TAM and UTAUT. Despite their widespread applications for examining the adoption of IT projects in the healthcare sector, a single theory or model may not consistently provide a sufficient explanation for the phenomena being investigated. Therefore, it is necessary to adopt a multifaceted approach to studying adoption. This can be achieved by using more than one model or theory and extending them by integrating additional contextual factors [29]. Such integration is essential for a better understanding of user technology acceptance, considering the intricacies of the IoB-based TM context from various viewpoints. Incorporating these factors allows for a more comprehensive and holistic understanding of user technology acceptance, which can differ based on the specific field context [30].

III. PROPOSED ADOPTION MODEL: IOB-TMAF

In the context of this study, the proposed model IoB-TMAF is developed based on a synthesis of systematic reviews on TM adoption literature conducted earlier [9]. This comprehensive model combines various well-established factors, offering a robust framework for studying the adoption of IoB-based TM, by providing valuable insights into different users' behaviour and intentions in this context. The TAM model is employed as the basis for the model, with its original four factors, namely Perceived Usefulness (PU), Perceived Ease of Use (PEoU), Attitude (ATT), and Behavioural Intention (BI). In addition, the model includes two factors from the UTAUT model (Social Influence (SI), and Facilitating Condition (FC)), one factor from HBM (Perceived Health Risk (PHR)), and other external variables, including Self-Efficacy (SE), Perceived Privacy Control (PPC), Perceived Cost (PC), Perceived Trust (PT), Task Fit (TF), and Reliability (R). The selected factors represent the most relevant and frequently identified predictors for this study's context, which examines individuals' intentions to adopt IoB-based TM [9]. These factors were chosen for their direct applicability to the study's goals. Furthermore, it is worth noting that TAM is frequently cited as the prevailing model for understanding acceptance in the healthcare domain due to its simplicity, flexibility, and ability to provide adequate explanatory power [31], [32].

In our proposed model, we estimated the effect sizes of these several predictors on IoB-based TM adoption intention. These estimates are derived from previous similar studies, converted to a common metric Cohen's f^2 , and then averaged across all studies to provide a comprehensive understanding of their impact. The f^2 values were interpreted according to Cohen's (1988) guidelines, where $0.02 \leq f^2 < 0.15$ represents a small effect, $0.15 \leq f^2 < 0.35$ represents a medium effect, and $f^2 \geq 0.35$ represents a large effect.

As shown in Table I, PU demonstrated the largest effect size ($f^2=0.782$), followed by PEoU ($f^2=0.529$) and ATT ($f^2=0.418$). These variables exhibited large effects on IoB-based TM adoption. TF also showed a large effect ($f^2=0.390$). Several variables demonstrated medium effects, including SI, SE, PT, FC, and PHR. R showed a small effect, while PPC and PC exhibited very small effects.

These findings suggest that interventions or strategies focusing on improving users' attitudes, perceived usefulness, and perceived ease of use may have the most substantial impact on IoB-based TM adoption. However, the influence of other factors should not be discounted, as even small effects can be meaningful in this context.

Fig. 1 shows the IoB-TMAF proposed model for the investigation. Based on the meta-analysis, the predictors of adopting IoB-based TM were categorized and organized into four main groups: individual context, technological context, health context, and organizational context. The following sections illustrate the interconnections between ideas and concepts related to the research problem and provide a detailed description of each category along with its respective factors.

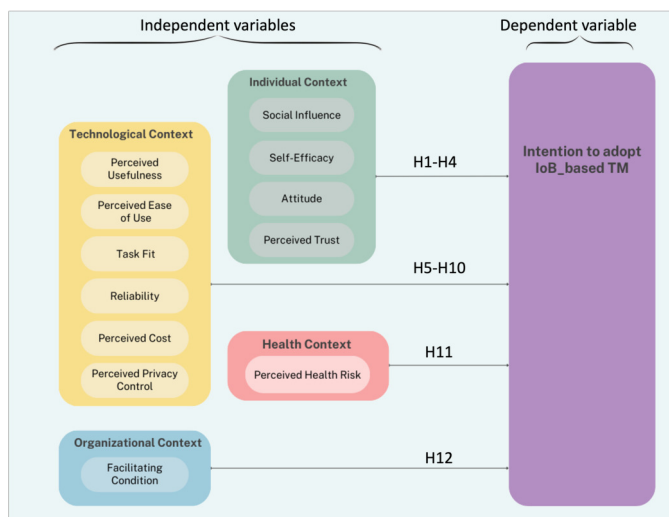


Fig. 1. IoB-TMAF Proposed Model.

A. Individual Context

Individual factors play a significant role in shaping a user's decision either to adopt or decline the usage of the technology [17]. Individual factors refer to those aspects that stem

from the individual themselves and influence their interaction with the system [44]. The individual context includes beliefs, thoughts, attitudes, and trust, which result from information and experiences, influencing decisions and shaping how individuals interpret various aspects. The review by [20] reported that individual resistance is considered one of the main barriers to TM implementation in Middle East countries. Examples of barriers related to individuals include a lack of awareness, knowledge, culture, trust, and motivation to use the technology [20]. Some researchers state that identifying and addressing potential barriers related to the individual context can facilitate a smoother transition to new technology and mitigate resistance [16], [20], [17]. Therefore, there is a need to consider the individual context to examine the key factors that determine the success of IoB-based TM adoption, providing more user-centered and adaptive technological solutions. The factors related to the individual context under investigation include attitude, social influence, self-efficacy, and trust beliefs, which are detailed below.

1) *Social Influence (SI)*: Numerous studies have highlighted the complicated role of SI in shaping the acceptance of emerging technologies [30], [45], [46]. SI, a key construct in the UTAUT model, is defined as "the degree to which an individual perceives that important other (e.g., family and friends) believe they should use the new system" [27]. In this study, SI refers to the extent to which an individual believes that their decision regarding the adoption of IoB devices in TM for diagnosis is influenced by the recommendations of others. Empirical evidence demonstrates that SI positively impacts an individual's intention to adopt various types of health information technology, including mobile health services [47], wearable health devices [30], [47], and TM [30], [33], [48], [34].

Building on these findings, this study assumes that SI plays a pivotal role in supporting individuals' confidence to embrace IoB-based TM and enriches their remote healthcare experience. Previous research underscores SI's significant impact on the adoption of TM by both healthcare providers [48], [34] and patients [30], [33]. A recent study by [49] supports the common observation that IoT consumers frequently seek advice from family members, peers, and colleagues uncertain about a product. Moreover, SI is driven by the influence of highly educated and successful individuals who share their experiences and motivating their social circles over time [33]. Therefore, the following hypothesis is proposed:

H1: Social Influence is significantly influencing user intention to adopt **IoB-based Telemedicine**.

2) *Self-Efficacy (SE)*: SE significantly influences individuals' attitudes in adopting and effectively utilizing new technological advancements, including various domains such as healthcare and technology-mediated applications. SE refers to an "individual's belief or judgment regarding their capability to use a technology to accomplish a particular job or task" [50]. The concept of SE is like Perceived Behavioural Control in the Cognitive Theory of the TPB [50], and Computer Self-Efficacy

TABLE I
ESTIMATE THE EFFECT SIZES OF PREDICTORS ON IOB-TMAF ADOPTION

Predictor variable	Estimated Effect Size f^2	Description	Source
SI	0.192	Medium effect	[33], [34], [35], [30], [36], [37], [38], [39], [40], [41]
SE	0.110	Medium effect	[33], [41], [17], [42], [36]
ATT	0.418	Large effect	[17], [32], [42], [35], [34], [39]
PT	0.116	Medium effect	[43], [42], [40], [37], [39]
PU	0.782	Large effect	[33], [32], [34], [41], [35], [30], [36], [37], [38], [39], [40]
PEoU	0.529	Large effect	[33], [32], [34], [41], [30], [36], [37], [38], [40], [17], [40]
TF	0.390	Large effect	[33], [39], [38]
R	0.087	Small effect	[39], [36]
PPC	0.033	Small effect	[39], [37]
PC	0.018	Small effect	[39], [30]
FC	0.185	Medium effect	[33], [30], [39], [34], [35], [17], [37], [38], [40]
PHR	0.171	Medium effect	[43], [36], [42], [30]

in the extended TAM [51]. It serves as internal control over the belief that an individual can carry out a specific behaviour [51] and as predictive in decision-making regarding technology use [52].

Prior research provides experimental evidence supporting the significant influence of SE on the intention to use new technology [33], [51], [53]. In IoB-based TM, SE represents an individual's belief in the set of skills they possess to use IoB devices in a TM setting. The literature in health information technology has shown a significant impact of SE on the adoption of TM applications for patients [47], [33], [54] and healthcare providers [17], [34]. Users with lower levels of SE are uncertain and less comfortable when using technology [50]. In a study conducted by [33], it was observed that patients who possessed a high sense of SE demonstrated a favorable intention toward the adoption of TM. This finding underscores the significant role that confidence in one's ability to perform specific tasks or actions in influencing individuals' receptiveness to technological innovations in healthcare.

Similarly, another study by [17] revealed that healthcare providers exhibited a greater willingness to adopt TM when they perceived themselves as competent and capable of effectively using the various devices and tools associated with TM services. This connection between SE and the willingness of healthcare professionals to engage with technology suggests that SE is a pivotal factor not only in the patient's acceptance of technology but also in the readiness of healthcare providers to integrate technological solutions into their practice. These findings collectively highlight the critical role of SE in shaping the adoption and effective utilization of technology within the healthcare industry. Therefore, the following hypothesis is proposed:

H2: Self-Efficacy is significantly influencing user intention to adopt **IoB-based Telemedicine**.

3) *Attitude (ATT)*: In the realm of technology adoption, each user brings a unique perspective and attitude toward a specific item, a notion emphasized by the TAM model [17]. TAM highlights that the construct of ATT shapes the individuals' intention to engage in certain behaviours related to technology [23]. ATT is defined as "the degree of evaluative effect that an individual associates with using the target

system in their job" [52]. The role of ATT stands out as a key predictor of behavioural intention in numerous studies [23], [55]. Interestingly, within the healthcare domain, ATT's significance diverges. While some studies suggest it may not significantly contribute to the intention to adopt technology [56], others identify it as a variable influencing the behavioural intentions of healthcare providers and patients, particularly in the context of technology acceptance such as TM [17], [47], [57]. This discrepancy underscores the complex link between individuals' sentiments regarding technology and their inclination to consider its use [23].

However, when a positive attitude toward technology is present, it can have significant benefits for both healthcare providers and patients. Researchers believe that the positive ATT of healthcare providers towards technology not only fosters increased commitment and enthusiasm for its adoption in their medical practice but also plays a significant role in improving patient care quality [32], [17]. This positive outlook among healthcare professionals encourages them to explore and implement cutting-edge technologies, which can lead to more efficient diagnoses, treatment options, and overall healthcare delivery [58].

Similarly, patients who have a favorable perception of the technology used in their healthcare experience greater confidence and trust in their medical providers [47]. This trust in the healthcare system's technological capabilities can result in higher patient satisfaction and better engagement in their own care [59]. Moreover, authors in [60] reported that when patients feel comfortable and empowered by the technology integrated into their healthcare journey, they are more likely to actively participate in managing their health and adhering to treatment plans, ultimately contributing to better health outcomes. Therefore, the following hypothesis is proposed:

H3: Attitude is significantly influencing user intention to adopt **IoB-based Telemedicine**.

4) *Perceived Trust (PT)*: Trust holds the power to shape a user's intentions and attitudes toward an innovative IT system [47]. This influence becomes even more crucial in healthcare, where an increasing number of services are being provided through technologies that need users' involvement, participation, and trust [47]. PT is defined as "faith in the

adoption of a new technology that end-users place in it with regards to the services this technology can provide" [43]. In the context of this study, PT, is defined as the degree of confidence and reliance that end-users place in the providers, IoB, and the TM system as a whole [47].

Several studies highlight the importance of trust in this context. Scholars argue that trust is a main precursor influencing the adoption of e-health [61]. It represents the most significant factor in fostering a successful healthcare relationship in remote doctor-patient communication [47]. Moreover, as the healthcare sector continues to integrate cutting-edge technologies, such as IoB, trust assumes an even more pivotal role [62]. A study by [63] emphasized the importance of trust for the IoT, as IoB falls under the IoT umbrella. The authors highlighted two key aspects: firstly, the interaction and trust among various IoT elements, including body-connected devices, and secondly, the trust of users in adopting and using these technologies [40], [63].

In addition to facilitating technology adoption, establishing trust is critical for the successful implementation of eHealth services. Studies have shown that an atmosphere of trust can enhance the efficacy of technology adoption [61], [40]. Conversely, a lack of trust in healthcare technology may adversely impact patients' health outcomes [47]. Similarly, distrust of IT has been identified as a key factor leading to the avoidance of using technology as a resource for health-related purposes [64]. Therefore, the impact of trust on user perceptions and behaviours towards IT systems, particularly in the healthcare sector, underscores its indispensable role in shaping the success of technological integration. Therefore, the following hypothesis is proposed:

H4: Perceived Trust is significantly influencing user intention to adopt **IoB-based Telemedicine**.

B. Technological Context

To build a successful ICT system, it is essential to not only establish the appropriate technological infrastructure but also develop a comprehensive understanding of user perceptions and behaviours from a technological perspective [65]. The technological concept encompasses the technological elements and strategic considerations involved in designing, implementing, establishing communication infrastructure, and utilizing related technologies to facilitate the acceptance of the new system- in our context, the provision of remote healthcare [9]. This includes, for example, system characteristics, system quality, complexity, security and privacy, information accuracy, cost, and job reflection [10].

Building on the insights of [65], [25], it is evident that explaining the technological aspects of an ICT system is pivotal in gaining user acceptance. The TAM model serves as a valuable guide, emphasizing that users' acceptance is influenced by the perceived ease of use and perceived usefulness of the new system [65]. To comprehend the adoption of a system, specifically the IoB-based TM in this study, it is imperative to identify key technological factors impacting users. The factors related to the technological context under investigation include

perceived usefulness, perceived ease of use, task fit, reliability, perceived cost, and perceived privacy control.

1) *Perceived Usefulness (PU)*: PU represents one of the core constructs of the TAM model [65]. It is defined as "the extent to which a person believes that using the system will enhance his or her job performance" [26]. In the context of IoB-based TM, PU can be viewed as the extent to which healthcare providers believe that using this technology will enhance their performance and productivity to provide effective and high-quality care to patients. For patients, it can be viewed as their belief that using this technology will improve their health management and overall healthcare experience. According to [66], IoB technology has the potential to transform healthcare by enabling continuous tracking of patients' vital signs and health information in real-time. This shift towards continuous remote patient monitoring can provide healthcare providers with the timely, accurate insights needed to make prompt diagnoses and treatment decisions.

Previous research has consistently identified PU as a key driver in the adoption of technology among healthcare providers [17], [56], [67] and patients [68]. The belief that integrating IoB into TM practices will enhance diagnostic capabilities serves as a crucial determinant for its acceptance. The study by [69] underscores that consumer adoption of healthcare wearable devices is significantly influenced by the perception that these devices enhance healthcare effectiveness, emphasizing the pivotal role of user beliefs in shaping technology adoption. Authors of [31] further expand the scope of PU in the healthcare domain, transcending individual productivity to encompass broader aspects such as increased efficiency, elevated quality and safety standards, enhanced workflow support, patient empowerment, and other utility measures specific to healthcare [70], [31].

Conversely, the absence of PU in information technology can pose a substantial barrier to adoption. As elucidated by [70] in their review, PU is not merely a singular barrier but a linchpin, constituting 15% of the various barriers encountered in the adoption of electronic health records in the KSA [49], [70]. Thus, understanding PU in the IoB-based TM context becomes important for fostering widespread adoption and ensuring successful implementation in healthcare settings in the KSA. Therefore, the following hypothesis is proposed:

H5: Perceived Usefulness is significantly influencing user intention to adopt **IoB-based Telemedicine**.

2) *Perceived Ease of Use (PEoU)*: In the TAM model [65], PEoU stands as one of the core constructs. It is defined as "the extent to which a person believes that using the system will be free of effort" [26]. In the context of this study, it refers to the user's perception of the ease of using IoB-based TM. The degree of ease associated with using technology is found to positively affect the acceptance behaviour [65], [27]. This indicates that the smoother and more straightforward the user perceives the system, the more likely they are to embrace and adopt it. This positive relationship between PEoU and user acceptance underscores the importance of designing and

implementing technology in a manner that minimizes cognitive load and operational complexities for the end user [65].

Providing support for these ideas, numerous studies have confirmed that PEOU significantly influences the adoption of technologies in the healthcare domain [21]. Authors of [47], [67] emphasized the paramount importance of user-friendliness and accessible technical support in telehealth systems. Additionally, [17] reported that healthcare provider providers exhibit a greater willingness to adopt and incorporate technology into their practices when they find it to be user-friendly and require minimal cognitive effort, given their complex working environment and busy schedules. This aligns with the broader literature on technology adoption, emphasizing that usability and ease of use are key determinants influencing users' willingness to engage with and embrace novel technological solutions [27], [51].

However, the complexity of technology, particularly in the healthcare domain, gives rise to resistance and rejection to adopt innovations [20], [70]. With specific reference to Saudi healthcare services, an identified challenge emerges in the form of a perceived lack of ease of use, which accounts for approximately 15% of the obstacles hindering the widespread adoption of electronic health records [70]. This implies the importance of this factor to facilitate the seamless integration of advanced technological solutions, such as IoB-based TM, in Saudi healthcare practices. Therefore, the following hypothesis is proposed:

H6: Perceived Ease of Use is significantly influencing user intention to adopt **IoB-based Telemedicine**.

3) *Task Fit (TF)*: In the realm of ICT systems, ensuring that the applications, and services align with the users' needs and objectives is critical. This involves conducting comprehensive needs assessments, understanding user requirements, and customizing the technological solutions to address specific challenges and goals [71]. TF is defined as "an individual's perception regarding the degree to which the target system is applicable to his or her job" [26]. In other words, it reflects how essential the IoB-based TM system's capability is to support a specific set of tasks related to diagnosis within that job. The TF concept is like the Job Relevance in the extended TAM [26], and Task-Technology Fit [72], which have empirically confirmed positive roles in healthcare acceptance of technology [33], [54].

Moreover, TF serves as a judgment that influences the individual's intention toward adopting technology [26]. According to the extended TAM model, individuals employ a mental representation to evaluate the alignment between significant work objectives and the outcomes of engaging with a system [26]. This serves as the foundation for making judgments regarding the perceived usefulness of that system [26]. Similarly, earlier studies in human-computer interaction and psychology have highlighted that users' specific knowledge, shaped by mental representations of their job situations, can function as a foundation for identifying tasks compatible with a given system [26], [73], [74].

In the context of TM, as reported by several studies, the lack of physical examination in TM is considered one of the main barriers in the field [35], [75], [76]. If the IoB provides this potential, healthcare providers and organizations are more likely to perceive the usefulness of the systems, as they are compatible with work needs and values. With IoB technologies, healthcare providers can remotely access real-time physiological data, allowing for a more accurate and dynamic evaluation of patients [58]. This capability not only addresses the current barrier of limited physical examination in TM but also opens new avenues for proactive and personalized healthcare. Therefore, the following hypothesis is proposed:

H7: Task Fit is significantly influencing user intention to adopt **IoB-based Telemedicine**.

4) *Reliability (R)*: The significance of R in system development becomes evident when assessing user adoption and utilization, emphasizing its pivotal role as a crucial quality factor [77]. Various studies delve into the R within user acceptance models, each identifying distinct dimensions based on the context and specific study objectives [78]. Common elements include system reliability, output quality, data accuracy, error handling, system availability, and response time. Previous research has highlighted that both system quality and information quality not only serve as key predictors for the adoption of IT but also exert a significant impact on user satisfaction [79].

Moving to the healthcare domain, [36] identified the health information technology (HIT) reliability, incorporating the quality of output and demonstrability of results. Output quality refers to "the degree to which an individual believes that the system performs his or her job tasks well" [25], [26], while result demonstrability refers to "the degree to which an individual believes that the results of using a system are tangible, observable, and communicable" [25], [26]. These concepts were derived from the extended TAM model, which posits that PU is shaped by both output quality and result demonstrability. Consequently, individuals assess the system's effectiveness and accuracy in task execution, forming more favorable perceptions of a system's usefulness when the correlation between usage and positive outcomes is easily observable [26].

In the context of this study, R is crucial to ensure that users trust and rely on the IoB devices for monitoring and assessing their health [66]. Several dimensions or factors contribute to the R of IoB-based TM systems in the context of user acceptance models. R, in this study, will be examined to measure contextual aspects, including output quality, data accuracy, and demonstrability of results, aligning with previous studies in this domain [69]. Data accuracy refers to the individual's perception of the accuracy of the information presented or processed by the system. Therefore, examining the R within user acceptance models, encompassing dimensions such as output quality, data accuracy, and demonstrability of results, underscores its indispensable role in fostering user trust, satisfaction, and reliance on IoB devices for health assessment

in the TM practice. Therefore, the following hypothesis is proposed:

H8: Reliability is significantly influencing user intention to adopt **IoB-based Telemedicine**.

5) *Perceived Cost (PC)*: Perceived cost is like the price value in the extended UTAUT model. It is defined as "the consumers' cognitive trade-off between the perceived benefits of the applications and the monetary cost for using them" [80]. Another study defined perceived cost as "concerns on the costs consumed in buying, using, and repairing the component of a particular system or service" [81]. Based on the study's context, it refers to the concerns about the cost, including various elements such as expenses related to data service operators (mobile Internet), device acquisition costs, and applicable service charges [30]. This multifaceted understanding of cost is crucial in evaluating the overall considerations that users weigh when adopting new technologies.

A body of prior research has consistently highlighted cost as a significant barrier to the utilization and acceptance of IoT products and services [30], [82]. Authors of [81] found that the cost is a significant determinant of the intention to adopt IoT technology within smart home environments. This suggests that understanding and managing the financial implications associated with adopting IoT technologies is essential for successful implementation and user acceptance. Similarly, [83] delved into the healthcare sector, emphasizing the relevance of the cost factor in doctors' intentions to use IoT healthcare devices, particularly during the challenging times of the COVID-19 pandemic. This underscores the importance of cost consideration as a hindering factor in the adoption process, influencing users' intentions as it is perceived as a hindrance in various contexts [81]. This highlights the role of the cost factor as a critical determinant in the acceptance and adoption of IoB technologies in a TM setting.

If users find IoB-based TM devices at an affordable cost, they are more likely to have a positive intention to adopt and accept them. Earlier studies have indicated that TM serves as a cost-effective means of accessing healthcare services remotely [84]. Consequently, researchers deduce that PC plays a significant role as a determinant of the behavioural intention to utilize the technology. The accessibility and convenience offered by TM have proven instrumental in overcoming geographical barriers and improving healthcare outcomes [15]. Therefore, the following hypothesis is proposed:

H9: Perceived Cost is significantly influencing user intention to adopt **IoB-based Telemedicine**.

6) *Perceived Privacy Control (PPC)*: In the realm of healthcare services, individuals' receptiveness to innovative technologies is significantly influenced by the concern about privacy [67]. The critical examination of privacy and data security issues remains pivotal in shaping their willingness to share information [85]. PPC is defined as "an individual's perception that healthcare providers have control over the amount of information to be shared and disclosed with others" [85]. To ensure accurate diagnoses and effective treatment, patients are

required to disclose their information to healthcare providers [68]. Nevertheless, the fear of social discrimination arises when it comes to sharing sensitive details like psychological and mental health issues, as well as conditions such as HIV, leading patients to hesitate in disclosing such information [68].

Given the sensitivity of personal health information for individuals, it becomes imperative to consider the impacts of data security factor when assessing an individual's acceptance of IoB-based TM devices [69]. Adoption of healthcare technology occurs when an individual perceives the benefits to outweigh the potential privacy risks, as indicated by [86]. Conversely, if this balance tips in favour of privacy risk, the technology is likely to be rejected [69], [86].

Regarding the adoption of TM in Saudi healthcare services, prior studies revealed that 90% of doctors expressed concerns regarding the privacy of patients [49], [86]. Interestingly, doctors are perceived to harbour the highest level of privacy concern, surpassing even that of the patients themselves [49]. This heightened level of concern among doctors may be attributed to their front-line role, which fosters an increased sense of responsibility for the protection of patient information.

Moreover, [67] revealed in their study that ensuring privacy in technological infrastructure requires a careful balance between confidentiality and practicality, as neglecting the latter can unintentionally hinder ease of use. This is particularly crucial in healthcare, where designing rational digital permissions is essential to safeguarding patient information while enabling efficient access for healthcare providers. Therefore, the following hypothesis is proposed:

H10: Perceived Privacy Control is significantly influencing user intention to adopt **IoB-based Telemedicine**.

C. Health Context

The impact of health considerations on individuals' perspectives, attitudes, and behaviours regarding the adoption of new healthcare technologies is significantly important [87]. Well-known factors include health interests, perceived health risks, perceived severity, perceived vulnerability, and health beliefs [28]. Understanding and addressing these considerations can help ensure that technologies are designed, implemented, and promoted to effectively meet user needs and concerns. Neglecting this could limit adoption or lead to suboptimal utilization. As healthcare technology evolves, understanding the complex interaction between health factors and technological advancements shapes how people perceive, embrace, or resist integration, influencing innovation and implementation [88], [87]. Notably for TM, studies show patients are more impacted by health factors than providers, indicating the significant role of patients' psychological and emotional considerations like concerns over remote consultation efficacy or misdiagnosis risk. Reviews highlight perceived health risk as pivotal in influencing technology acceptance in healthcare, underscoring its importance for frameworks like IoB-TMAF [9], [10].

Perceived Health Risk (PHR): The significance of risk as a pivotal determinant of human behaviour cannot be overlooked.

Perceived risk refers to an individual's assessment of the risk when deciding to engage in a particular action or activity [30]. The effects of risk and uncertainty are unavoidable in the realms of health and information communication technologies. The HBM theorizes that a person's health-related beliefs influence their health behaviours based on perceived susceptibility, perceived severity, perceived benefits of action, perceived barriers to action, and self-efficacy [28]. Additionally, the study [37] has classified perceived barriers in healthcare into seven groups, namely, time, financial, performance, privacy, physiological, social, and overall risk.

Building on that, this research study defines PHR for patients in terms of susceptibility, severity, and performance risk. The other factors were omitted either as they are described in the framework as separate factors or are based on the organizational level, which is out of the scope of this study. Performance Risk refers to the probabilistic perception that a TM system may harm patients by failing to provide adequate information about their health status due to the lack of physical examination [37]. Perceived Susceptibility refers to "an individual's belief about their likelihood of getting a health condition or problem" [28]. Perceived Severity is defined as "an individual's belief about the seriousness of a health condition or its potential consequences" [28]. Therefore, integrating the IoB with TM may shape individuals' beliefs and intentions by reducing these kinds of risks and enhancing their interest in health to engage in such practices.

According to the HBM, an individual's likelihood of engaging in health-related behaviours depends on their perceived susceptibility to and severity of a health threat [28]. When people believe they are more vulnerable to a condition or that it would have very serious personal consequences, they become more motivated to act to prevent its onset [89]. Applying this model to the adoption of online health services, research shows that higher perceived susceptibility to and severity of health threats are associated with stronger intentions to use these technologies as part of one's health management, as observed in the study by [53]. Essentially, the more an online health information seeker believes they are at risk for and could be seriously impacted by a health issue if they do not act, the more likely they are to adopt available online health services to empower their health decisions and behaviours.

Furthermore, research shows that when people perceive a health threat to be more severe or likely, they are more likely to adopt new health technologies to reduce and mitigate that threat [53], [36]. These studies found a positive relationship between an individual's risk assessment of a health issue and their intention to adopt health-related technology. Thus, understanding and addressing the multifaceted dimensions of perceived risk, including perceived threat, and performance aspects, is crucial for shaping individuals' beliefs and intentions, and fostering greater interest and engagement in health-related technologies. Therefore, the following hypothesis is proposed:

H11: Perceived Health Risk is significantly influencing user intention to adopt **IoB-based Telemedicine**.

D. Organizational Context

The organizational context plays a crucial role in explaining user acceptance of new technology. It refers to the various internal elements within an organization that influence the user's intention regarding the extent to which new technologies are embraced, implemented, and integrated into their operations [9]. These factors encompass a wide range of aspects related to the organization's culture, top management support, resources availability, and the alignment with existing system [20]. The literature review has clearly demonstrated that the control of the external factors, i.e., the support and resources accessible to the individual to enable them to engage in the behaviour, are the major influencers for adopting new technologies [33], [52]. The increased likelihood of investing in new technology is underscored by [17], who highlight that organizations with greater top management support and superior IT capabilities are more likely to adopt such advancements. Further, in terms of IoB-based TM, a recent review identified that the FC is the most significant factor under organizational context, influencing the adoption of the system for both patients and providers [9]. Thus, the factor related to the organizational context under investigation include FC represented in resources and management supports.

Facilitating Condition (FC): FC plays a paramount role in shaping the environment for technological adoption within an organization. These conditions encompass a set of elements ranging from adequate financial resources to the provision of skilled manpower, training courses, and technical infrastructure [46]. FC, one of the core constructs in the UTAUT model, is defined as "the degree to which an individual believes that an organizational and technical infrastructure exists to support the use of the system" [27]. Within the scope of this study, FC refers to an individual's perception of possessing the requisite knowledge, resources, and support for engaging in the IoB-based TM system.

The tendency to adopt technology in the TM is significantly influenced by sufficient technical infrastructure and organizational support [33]. As highlighted by [35], who found that healthcare providers are more motivated to engage in teleneurology when supported by robust infrastructure and organizational backing. These positive effects tend to amplify over time due to the availability of continuous assistance and guidance, underscoring the critical role of robust infrastructure and resources in sustaining technological integration [17], [35]. Additionally, top management support is crucial, as leadership commitment to innovation fosters an organizational culture open to change and experimentation, enhancing the acceptance of new technology [17]. Studies show that top management's commitment positively impacts the reception of new systems, promoting beliefs in their usefulness and ease of use, thereby facilitating widespread acceptance and successful integration within the organization [17]. Therefore, the following hypothesis is proposed:

H12: Facilitating Condition is significantly influencing user intention to adopt **IoB-based Telemedicine**.

IV. CONCLUSION AND FUTURE WORK

This study aimed to investigate the core factors influencing the acceptance of adopting IoB-based TM among Saudi patients and healthcare providers. The proposed IoB-TMAF framework, grounded in the TAM, was developed to identify the key factors influencing the acceptance of this system by its primary users. These factors are derived from a variety of contexts including individual, technology, organization, and health.

The proposed model emerged as an important finding from the systematic review and the initial year of thesis work. It is expected to serve as a valuable resource for the Saudi Ministry of Health (MOH), healthcare policymakers, and practitioners by providing critical factors for the successful utilization of IoB-based TM in supporting diagnosis from the perspective of end-users, including providers and patients. Consequently, this is anticipated to improve the quality and efficiency of health services.

While this study lays a foundational framework, it is limited by its theoretical nature at this stage. The next steps involve empirical validation which is crucial for confirming the framework's applicability in real-world settings. Future research will involve a comprehensive study to gather data and analyze the model's performance and validity using a mixed method. This will include two sequential phases: a qualitative data collection phase involving semi-structured interviews with a small sample of twelve healthcare providers and patients, and a quantitative data collection phase through an electronic survey targeting at least 450 participants. The qualitative phase aims to explore the proposed model and other construct-based theories of behaviour by identifying users' feelings and perceptions regarding the implementation of IoB-based TM, allowing for potential refinements to the model. In the quantitative phase, the hypotheses will be tested, the proposed model will be empirically examined using Structural Equation Modeling (SEM), and the estimated effect size will be validated and refined.

By continuing this research, we aim to provide robust, empirically validated insights into the factors influencing the acceptance of IoB-based TM. This research has the potential to significantly enhance the quality and efficiency of healthcare services in Saudi Arabia by providing a model that is responsive to the needs and perspectives of both healthcare providers and patients.

REFERENCES

- [1] I. Grand View Research. (2023, Dec) Telemedicine market size to reach \$380.3 billion by 2030. [Online]. Available: "https://www.grandviewresearch.com/press-release/global-telemedicine-industry"
- [2] A. Albahri, J. Alwan, Z. Taha, S. Fawzi, R. Amjed, A. Zaidan, O. Albahri, B. Bahaa, A. Alamoodi, and M. Alsalem, "Iot-based telemedicine for disease prevention and health promotion: State-of-the-art," *Journal of Network and Computer Applications*, 10 2020.
- [3] F. Albejaidi and K. Nair, "Building the health workforce: Saudi arabia's challenges in achieving vision 2030," *The International Journal of Health Planning and Management*, vol. 34, 08 2019.
- [4] M. Mohaya, M. Almaziad, K. Al-Hamad, and M. Mustafa, "Telemedicine among oral medicine practitioners during covid-19 pandemic and its future impact on the specialty," *Risk Management and Healthcare Policy*, vol. Volume 14, pp. 4369–4378, 10 2021.
- [5] S. Sood, V. Mbarika, S. Jugoo, R. Dookhy, C. Doarn, N. Prakash, and R. Merrell, "What is telemedicine? a collection of 104 peer-reviewed perspectives and theoretical underpinnings," *Telemedicine journal and e-health : the official journal of the American Telemedicine Association*, vol. 13, pp. 573–90, 11 2007.
- [6] M. Ong, P. Romano, S. Edgington, H. Aronow, A. Auerbach, J. Black, T. Marco, J. Escarce, L. Evangelista, B. Hanna, T. Ganiats, B. Greenberg, S. Greenfield, S. Kaplan, A. Kimchi, H. Liu, D. Lombardo, C. Mangione, B. Sadeghi, and G. Fonarow, "Effectiveness of remote patient monitoring after discharge of hospitalized patients with heart failure," *JAMA Internal Medicine*, vol. 176, 02 2016.
- [7] A. Di Cerbo, J. C. Morales Medina, B. Palmieri, and T. Iannitti, "Narrative review of telemedicine consultation in medical practice," *Patient Preference and Adherence*, vol. 9, pp. 65–75, 01 2015.
- [8] A. Neamah, M. Ghani, A. Ahmad, E. Alomari, and R. R. Nuiaa, "E-health state in middle east countries: An overview," pp. 2974–2990, 09 2018.
- [9] T. Ghiwaa, I. Khan, M. White, and N. Beloff, "Telemedicine adoption for healthcare delivery: A systematic review," *International Journal of Advanced Computer Science and Applications*, vol. 14, 01 2023.
- [10] M. Al-rawashdeh, P. Keikhosrokiani, B. Belaton, M. Alawida, and A. Zwiri, "Iot adoption and application for smart healthcare: A systematic review," *Sensors*, vol. 22, p. 5377, 07 2022.
- [11] C. Amato, "Internet of bodies: Digital content directive, and beyond," *JIPITEC – Journal of Intellectual Property, Information Technology and E-Commerce Law*, vol. 12, 2021.
- [12] M. Zayyad and M. Toycan, "Factors affecting sustainable adoption of e-health technology in developing countries: An exploratory survey of nigerian hospitals from the perspective of healthcare professionals," *PeerJ*, vol. 6, p. e4436, 03 2018.
- [13] F. Alanezi, "Factors affecting the adoption of e-health system in the kingdom of saudi arabia," *International Health*, vol. 13, 11 2020.
- [14] W. H. O. (WHO). (1998, Dec) A health telematics policy in support of who's health-for-all strategy for global health development. [Online]. Available: https://apps.who.int/gb/ebwha/pdf_files/EB101/pdfangl/angid9.pdf
- [15] R. Wootton, "Telemedicine support for the developing world," *Journal of telemedicine and telecare*, vol. 14, pp. 109–14, 02 2008.
- [16] A. Alaboudi, A. Atkins, B. Sharp, A. Balkhair, M. Alzahrani, and T. Sunbul, "Barriers and challenges in adopting saudi telemedicine network: The perceptions of decision makers of healthcare facilities in saudi arabia," *Journal of Infection and Public Health*, vol. 9, 09 2016.
- [17] S. Zailani, M. Gilani, D. Nikbin, and M. Iranmanesh, "Determinants of telemedicine acceptance in selected public hospitals in malaysia: Clinical perspective," *Journal of medical systems*, vol. 38, p. 111, 09 2014.
- [18] M. Douglas, J. Xu, A. Heggs, G. Wrenn, D. Mack, and G. Rust, "Assessing telemedicine utilization by using medicaid claims data," *Psychiatric Services*, vol. 68, pp. 173–178, 02 2017.
- [19] H. Nadri, B. Rahimi, H. Lotfnezhad Afshar, M. Samadbeik, and A. Garavand, "Factors affecting acceptance of hospital information systems based on extended technology acceptance model: A case study in three paraclinical departments," *Appl Clin Inform*, vol. 09, pp. 238–247, 04 2018.
- [20] H. Al-Samarraie, S. Ghazal, A. Alzahrani, and L. Moody, "Telemedicine in middle eastern countries: Progress, barriers, and policy recommendations," *International Journal of Medical Informatics*, vol. 141, p. 104232, 07 2020.
- [21] A. Alqudah, M. Al-Emran, and K. Shaalan, "Technology acceptance in healthcare: A systematic review," *Applied Sciences*, vol. 11, p. 10537, 11 2021.
- [22] E. M. Rogers, *Diffusion of Innovations, 4th Edition*. New York: the Free Press, 1995.
- [23] M. Fishbein, I. Ajzen, and A. Belief, "Belief, attitude, intention, and behavior: An introduction to theory and research," *Contemporary Sociology*, vol. 6, 03 1977.
- [24] I. Ajzen, "The theory of planned behaviour: Reactions and reflections," *Psychology & health*, vol. 26, pp. 1113–27, 09 2011.
- [25] V. Venkatesh and H. Bala, "Technology acceptance model 3 and a research agenda on interventions," *Decision Sciences - DECISION SCI*, vol. 39, pp. 273–315, 05 2008.

- [26] V. Venkatesh and F. Davis, "A theoretical extension of the technology acceptance model: Four longitudinal field studies," *Management Science*, vol. 46, pp. 186–204, 02 2000.
- [27] V. Venkatesh, M. Morris, G. Davis, and F. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quarterly*, vol. 27, pp. 425–478, 09 2003.
- [28] I. M. Rosenstock, "Historical origins of the health belief model," *Health Education Monographs*, vol. 2, no. 4, pp. 328–335, 1974. [Online]. Available: <https://doi.org/10.1177/109019817400200403>
- [29] C. Barrette, "Usefulness of technology adoption research in introducing an online workbook," *System*, vol. 49, 04 2015.
- [30] W. Ben Arfi, I. Nasr, T. Khvatova, and Y. Ben Zaied, "Understanding acceptance of ehealthcare by iot natives and iot immigrants: An integrated model of utaut, perceived risk, and financial cost," *Technological Forecasting and Social Change*, vol. 163, 11 2020.
- [31] R. Holden and B.-T. Karsh, "The technology acceptance model: Its past and its future in health care," *Journal of biomedical informatics*, vol. 43, pp. 159–72, 08 2009.
- [32] J. Monthuy-Blanc, S. Bouchard, C. Mañano, and M. Seguin, "Factors influencing mental health providers' intention to use telepsychotherapy in first nations communities," *Transcultural psychiatry*, vol. 50, 05 2013.
- [33] M. Yamin and B. Alyoubi, "Adoption of telemedicine applications among saudi citizens during covid-19 pandemic: An alternative health delivery system," *Journal of Infection and Public Health*, vol. 13, 10 2020.
- [34] R. Evering, M. Postel, H. van Os-Medendorp, M. Bults, and M. den Ouden, "Intention of healthcare providers to use video-communication in terminal care: a cross-sectional study," *BMC Palliative Care*, vol. 21, 11 2022.
- [35] G. Pagaling, A. Espiritu, M. Dellosa, C. F. Leochico, and P. Pasco, "The practice of teleneurology in the philippines during the covid-19 pandemic," *Neurological Sciences*, vol. 43, 11 2021.
- [36] J. Kim and H.-A. Park, "Development of a health information technology acceptance model using consumers' health behavior intention," *Journal of medical Internet research*, vol. 14, p. e133, 10 2012.
- [37] S. Kamal, M. Shafiq, and P. Kakria, "Investigating acceptance of telemedicine services through an extended technology acceptance model (tam)," *Technology in Society*, vol. 60, p. 101212, 11 2019.
- [38] H. Wang, N. Yu, and X. Qu, "Understanding consumer acceptance of healthcare wearable devices: An integrated model of utaut and ttf," *International Journal of Medical Informatics*, vol. 139, p. 104156, 04 2020.
- [39] F. Dany and B. Römer, "Understanding dr. no - a comprehensive model explaining physicians' acceptance of telemedical systems," *ECIS 2014 Proceedings - 22nd European Conference on Information Systems*, 01 2014.
- [40] W. Ben Arfi, I. Nasr, G. Kondrateva, and L. Hikkerova, "The role of trust in intention to use the iot in ehealth: Application of the modified utaut in a consumer context," *Technological Forecasting and Social Change*, vol. 167, p. 120688, 06 2021.
- [41] A. Esber, M. Teufel, L. Jahre, J. Schmitt, E.-M. Skoda, and A. Bäuerle, "Predictors of patients' acceptance of video consultation in general practice during the coronavirus disease 2019 pandemic applying the unified theory of acceptance and use of technology model," *DIGITAL HEALTH*, vol. 9, p. 205520762211493, 01 2023.
- [42] M. I. Hossain, A. Fadhil, A. Hussin, N. Iahad, and A. Sadiq, "Factors influencing adoption model of continuous glucose monitoring devices for internet of things healthcare," *Internet of Things*, vol. 15, p. 100353, 01 2021.
- [43] N. Dogra, S. Bakshi, and A. Gupta, "Exploring the switching intention of patients to e-health consultations platforms: blending inertia with push-pull-mooring framework," *Journal of Asia Business Studies*, vol. ahead-of-print, 01 2022.
- [44] S. Kaphle, S. Chaturvedi, I. Chaudhuri, R. Krishnan, and N. Lesh, "Adoption and usage of mhealth technology on quality and experience of care provided by frontline workers: Observations from rural india," *JMIR mHealth and uHealth*, vol. 3, p. e61, 05 2015.
- [45] M. Mital, P. Choudhary, V. Chang, A. Papa, and A. Pani, "Adoption of internet of things in india: A test of competing models using a structured equation modeling approach," *Technological Forecasting and Social Change*, vol. 136, 03 2017.
- [46] A. Bhattacharjee and N. Hikmet, "Reconceptualizing organizational support and its effect on information technology usage: Evidence from the health care sector," *Journal of Computer Information Systems*, vol. 48, pp. 69–76, 06 2008.
- [47] J. Liu, G. Sorwar, M. Rahman, and M. Hoque, "The role of trust and habit in the adoption of mhealth by older adults in hong kong: a healthcare technology service acceptance (htsa) model," *BMC geriatrics*, vol. 23, p. 73, 02 2023.
- [48] C. Bakker, J. Huirne, F. Schaafsma, C. Geus, H. Bonjer, and J. Anema, "Electronic health program to empower patients in returning to normal activities after colorectal surgical procedures: Mixed-methods process evaluation alongside a randomized controlled trial (preprint)," 04 2018.
- [49] A. Alomari and B. Soh, "Determinants of medical internet of things adoption in healthcare and the role of demographic factors incorporating modified utaut," *International Journal of Advanced Computer Science and Applications*, vol. 14, 01 2023.
- [50] D. Compeau and C. Higgins, "Computer self-efficacy: Development of a measure and initial test," *MIS Quarterly*, vol. 19, pp. 189–211, 06 1995.
- [51] V. Venkatesh, "Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model," *Information Systems Research*, vol. 11, pp. 342–365, 12 2000.
- [52] S. RMarikyan, D. Papagiannidis, *THEORYHUB BOOK, A THEORY RESOURCE FOR STUDENTS AND RESEARCHERS ALIKE*. Newcastle upon Tyne, 2022. [Online]. Available: <https://open.ncl.ac.uk/>
- [53] Y. Sun, N. Wang, X. Guo, and J. Peng, "Understanding the acceptance of mobile health services: A comparison and integration of alternative models," *Journal of Electronic Commerce Research*, vol. 14, pp. 183–200, 01 2013.
- [54] H. Cho, G. Sanabria, M. Saylor, M. Gradilla, and R. Schnell, "Use of the fit framework to understand patients' experiences using a real-time medication monitoring pill bottle linked to a mobile-based hiv self-management app: A qualitative study," *International Journal of Medical Informatics*, vol. 131, 08 2019.
- [55] L. Russo, I. Campagna, B. Ferretti, E. Agricola, E. Pandolfi, E. Carloni, A. D'Ambrosio, F. Gesualdo, and A. Tozzi, "What drives attitude towards telemedicine among families of pediatric patients? a survey," *BMC Pediatrics*, vol. 17, 01 2017.
- [56] V. Brunelli, J. Fox, and D. Langbecker, "Disparity in cancer survivorship care: A cross-sectional study of telehealth use among cancer nurses in australia," *Collegian*, vol. 28, 01 2021.
- [57] P. Chau and P. Hu, "Investigating health care professionals decisions to accept telemedicine technology: An empirical test of competing theories," *Information & Management*, vol. 39, pp. 297–311, 01 2002.
- [58] S. Chakraborty and V. Bhatt, "Mobile iot adoption as antecedent to care-service efficiency and improvement: Empirical study in healthcare-context," *Journal of International Technology and Information Management*, vol. 28, p. 2019, 12 2019.
- [59] L. Liu and L. Shi, "Chinese patients' intention to use different types of internet hospitals: Cross-sectional study on virtual visits," *Journal of Medical Internet Research*, vol. 23, p. e25978, 08 2021.
- [60] H. Hoas, H. Andreassen, L. Lien, A. Hjalmarsen, and P. Zanaboni, "Adherence and factors affecting satisfaction in long-term telerehabilitation for patients with chronic obstructive pulmonary disease: a mixed methods study," *BMC Medical Informatics and Decision Making*, vol. 16, 02 2016.
- [61] W. Luo and M. Najdawi, "Trust-building measures: A review of consumer health portals," *Commun. ACM*, vol. 47, pp. 108–113, 01 2004.
- [62] D. McKnight, M. Carter, J. Thatcher, and P. Clay, "Trust in a specific technology: An investigation of its components and measures," *ACM Trans. Management Inf. Syst.*, vol. 2, p. 12, 01 2011.
- [63] J. Nord, A. Koohang, and J. Paliszkievicz, "The internet of things: Review and theoretical framework," *Expert Systems with Applications*, vol. 133, 05 2019.
- [64] D. Zulman, M. Kirch, K. Zheng, and L. An, "Trust in the internet as a health resource among older adults: Analysis of data from a nationally representative survey," *Journal of medical Internet research*, vol. 13, p. e19, 02 2011.
- [65] F. Davis and F. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quarterly*, vol. 13, pp. 319–, 09 1989.
- [66] R. Saravanakumar, P. Bedi, O. Hemakesavulu, T. Natarajan, E. Poornima, L. Thangavelu, and D. Jayadevappa, "Tob: Sensors for wearable monitoring and enhancing health care systems," *IEEE Instrumentation & Measurement Magazine*, vol. 25, pp. 63–70, 05 2022.

- [67] M. Nguyen, J. Fujioka, K. Wentlandt, N. Onabajo, I. Wong, R. Bhatia, O. Bhattacharyya, and V. Stamenova, "Using the technology acceptance model to explore health provider and administrator perceptions of the usefulness and ease of using technology in palliative care," *BMC palliative care*, vol. 19, p. 138, 09 2020.
- [68] S. Marhefka, D. Turner, and E. Lockhart, "Understanding women's willingness to use e-health for hiv-related services: A novel application of the technology readiness and acceptance model to a highly stigmatized medical condition," *Telemedicine and e-Health*, vol. 25, 08 2018.
- [69] Y. Gao, H. Li, and Y. Luo, "An empirical study of wearable technology acceptance in healthcare," *Industrial Management & Data Systems*, vol. 115, pp. 1704–1723, 10 2015.
- [70] A. Asma, R. Crowder, and G. Wills, "Barriers to the adoption of ehr systems in the kingdom of saudi arabia: an exploratory study using a systematic literature review," *Journal of Health Informatics in Developing Countries*, vol. 11, 07 2017.
- [71] A. Kakar, "A user-centric typology of information system requirements," *Journal of Organizational and End User Computing*, vol. 28, pp. 32–55, 01 2016.
- [72] D. Goodhue, "Understanding user evaluations of information systems," *Management Science*, vol. 41, pp. 1827–1844, 12 1995.
- [73] E. Locke and G. Latham, "A theory of goal setting & task performance," *The Academy of Management Review*, vol. 16, 04 1991.
- [74] D. Kieras and P. Polson, "Approach to the formal analysis of user complexity," *International Journal of Man-Machine Studies*, vol. 51, pp. 365–394, 04 1985.
- [75] A. Mubarak, A. Alrabie, A. Sibyani, R. Aljuaid, A. Bajaber, and M. Mubarak, "Advantages and disadvantages of telemedicine during the covid-19 pandemic era among physicians in taif, saudi arabia," *Saudi Medical Journal*, vol. 42, pp. 110–115, 01 2021.
- [76] M. Alsaleh, V. Watzlaf, D. DeAlmeida, and A. Saptono, "Evaluation of a telehealth application (sehha) used during the covid-19 pandemic in saudi arabia: Provider experience and satisfaction," *Perspectives in health information management*, vol. 18, p. 1b, 10 2021.
- [77] W. Wilkowska and M. Ziefle, "Understanding trust in medical technologies," 01 2018, pp. 62–73.
- [78] T. Hess, A. McNab, and K. Basoglu, "Reliability generalization of perceived ease of use, perceived usefulness, and behavioral intentions," *MIS Quarterly*, vol. 38, p. 2014, 03 2014.
- [79] W. Delone and E. McLean, "Information systems success: The quest for the dependent variable," *Information Systems Research*, vol. 3, pp. 60–95, 03 1992.
- [80] V. Venkatesh, J. Thong, and X. Xu, "Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology," *MIS Quarterly*, vol. 36, pp. 157–178, 03 2012.
- [81] E. Park, Y. Cho, J. Han, and S. Kwon, "Comprehensive approaches to user acceptance of internet of things in a smart home environment," *IEEE Internet of Things Journal*, vol. PP, pp. 1–1, 09 2017.
- [82] J. Nagy, J. Oláh, E. Erdei, D. Máté, and J. Popp, "The role and impact of industry 4.0 and the internet of things on the business strategy of the value chain—the case of hungary," *Sustainability*, vol. 10, p. 3491, 09 2018.
- [83] A. Alhasan, L. Audah, I. Ibrahim, A. Alsharaa, A. Al-Ogaili, and M. Jabiry, "A case-study to examine doctors' intentions to use iot healthcare devices in iraq during covid-19 pandemic," *International Journal of Pervasive Computing and Communications*, vol. ahead-of-print, 11 2020.
- [84] M. Abdellatif and W. Mohamed, "Telemedicine: An iot based remote healthcare system," *International Journal of Online and Biomedical Engineering (iJOE)*, vol. 16, p. 72, 05 2020.
- [85] X. Zhang, S. Liu, X. Chen, L. Wang, B. Gao, and Q. Zhu, "Health information privacy concerns, antecedents, and information disclosure intention in online health communities," *Information & Management*, vol. 55, 05 2018.
- [86] H. Li, A. Gupta, J. Zhang, and R. Sarathy, "Examining the decision to use standalone personal health record systems as a trust-enabled fair social contact," *Decision Support Systems*, vol. 57, p. 376–386, 01 2014.
- [87] S. Mustafa, W. Zhang, M. Shehzad, A. Anwar, and G. Rubakula, "Does health consciousness matter to adopt new technology? an integrated model of utaut2 with sem-fsqca approach," *Frontiers in Psychology*, vol. 13, 02 2022.
- [88] N. Abolhassani, B. Santos-Eggimann, A. Chiolero, V. Santschi, and Y. Henchoz, "Readiness to accept health information and communication technologies: A population-based survey of community-dwelling older adults," *International Journal of Medical Informatics*, vol. 130, 08 2019.
- [89] J. Mou and J. Cohen, "Trust, risk barriers and health beliefs in consumer acceptance of online health services," 01 2014.

Linked Labor Market Data: Towards a novel data housing strategy

Kristine Hein*

* Federal Institute for Vocational Education and Training (BIBB), Bonn, Germany

Abstract—The labor market is a domain rich in diverse data structures, both quantitative and qualitative, and numerous applications. This leads to challenges in the domain of data warehouse architecture and linked data. In this context, only a few approaches exist to generate linked data sets. For example, the multilingual classification system of European Skills, Competences, Qualifications, and Occupations (ESCO) and the German Labor Market Ontology (GLMO) serve as prominent examples showcasing the pivotal role of ontologies.

This paper introduces an initial conceptualization and proof-of-concept for managing interoperable German labor market data, including qualitative and quantitative data, such as surveys and statistical data, as well as textual data, such as social media data or online job advertisements. Additionally, it presents a data management perspective on the research network infrastructure, with a particular focus on the challenges encountered when establishing a data warehouse architecture within the field of education management. In this context, vocational training research offers a unique opportunity to anticipate future developments in the education and training markets. To this end, however, a fast and qualitatively good analysis option must be created to meet the demands of our fast-paced modern world. This is why a novel automated data strategy is required to facilitate the accelerated automation of processes, including ETL and the utilisation of contemporary data stacks.

I. INTRODUCTION

LABOR markets are dynamic entities, shaped by political and technical innovations, as well as changes in society. These shifts result in new demands, requirements, and novel data. In this context, vocational education and training, as well as re-training, play a pivotal role in meeting these new demands [1], [2]. However, labor market research is also an area with a limited amount of available resources [3]. It encompasses a multitude of data structures, each serving a distinct purpose. These include the facilitation of connections between job seekers and the most suitable training or employment opportunities. However, the sole application of semantic web technologies in this domain is the multilingual classification of European Skills, Competences, Qualifications and Occupations (ESCO). This exemplifies the pivotal role of ontologies in this field, see [4], [5].

We will now proceed to discuss the manifold problems and the problem statements. Following this, we will discuss the research questions of this work.

A. Problem statement

The sheer volume of data is overwhelming. While data is collected by all, there is no standardised structure for its collection, format or preparation, let alone documentation.

Gaining an overview of the individual data collections is challenging, and it is important to first cluster them and understand them at different levels. The field of research encompasses a variety of data types, which, in the context of survey data, can be broadly classified into three categories: survey data, process data, and structural data.

In the context of survey data, it is important to distinguish between qualitative and quantitative data. In the field of data science, the format of the data is initially distinguished from its content. This distinction is made, for instance, between text, image, video data, or formats such as XML and JSON. Furthermore, there is the realm of relational data (databases) in which standardised formats can be searched using a variety of search tools. These formats are designed with the objective of facilitating rapid searches and indexing.

Examples of data include survey data such as employment surveys, economic surveys such as the microcensus, job advertisement data, as well as already aggregated and harmonised data from economic structural research. Furthermore, data from social media, digitised archive data, structured data from education and training portals, market research data, and other sources are also utilised. Furthermore, qualitative data from internal BIBB surveys is also available. The data can be observed over time. The data is collected and updated at regular intervals, ranging from daily to once every six to ten years.

In the context of vocational training data, the occupation represents the common denominator of all data sets. However, some datasets can only be linked to each other via sectors, as surveys of companies (as a survey unit) and not of individuals were conducted.

Another challenge that arises in the context of surveys is the data protection of personal data. The person must have given their consent for the processing of GDPR-relevant data. As this consent is earmarked for a specific purpose, in some cases it is not possible to link the raw data at all. However, possibilities of abstraction and anonymisation can help here in order to still be able to work with the data collected. As a rule, the data volumes can be (artificially) increased by clever clustering so that conclusions cannot be drawn. Further work can take the direction of data boosting (see bootstrapping procedure [6]). In addition to these two defining characteristics, other common parameters of the datasets can be identified in certain instances, such as mapping to regions such as federal states, federal regions, cities/municipalities, districts, and so forth.

In addition to the amount of different data, there are also very different stakeholders at different levels of the data

warehouse. Standardisation of the input data. Uniqueness of the data records, storage in accordance with FAIR Data, manipulation- and access-protected. Ensure a minimalist principle for forwarding the data, and yet disclose it sufficiently so that the full scope of the data can be recognised, so that scientists have a comprehensive insight into and overview of the data.

Another challenge is the necessity of dealing with detailed data in special occupational areas or data gaps in other data sets. In such instances, estimates may be employed, provided that they are documented and labelled. In summary, the identification of ‘fuzzy’ data and estimates is to be achieved through the implementation of different data modelling techniques, which must be documented and disclosed in accordance with the relevant data pipelines. This is intended to ensure the traceability and reproducibility of the data by other scientists.

B. Research questions

In order to address the multitude of issues that have been identified, we will commence with three preliminary research inquiries.

- 1) What are the most effective methods for integrating labor market data into a data warehouse system?
- 2) What difficulties are encountered when conducting quantitative labor market research?
- 3) What specific challenges arise in the analysis of labor market research data, particularly in the context of historical, qualitative, and quantitative data?

The initial research question is relatively broad in scope, whereas the subsequent inquiries are more specific to the field of computational social sciences, with a particular focus on labor market research.

This paper is divided into six sections. The introductory section provides an overview of the subject matter and its relevance to the field. The second section offers a concise analysis of the current state of the art and related work. The third section delineates the methodological approach employed in this study. The fourth section presents the results and an evaluation of the approach. The final section presents the conclusions and offers a prospective outlook.

II. BACKGROUND AND LITERATURE REVIEW

Over the past decades, we find a growing interest in mining data from labor market data, educational databases, and information systems, see for example [7], [8], [9], [10], [3]. These studies have highlighted the importance of supporting decision-making and process management in labor market research. Data warehousing is a frequently employed methodology in the computational social sciences and big data pipelines [11], [12]. The generic challenges are typically the automated extraction of knowledge from data, which is usually interpreted passages from texts, and the mapping to existing data sets. However, there are still several challenges related to data and data integration. The research questions addressed in this field of study are diverse, encompassing topics such as occupational inequality [13], [14], migration and

language skills [15], sustainability [16], discrimination [17], and students and later occupation [18].

The situation in German-speaking countries (Germany, Austria, and Switzerland) with regard to automated analysis of labor market data is not significantly different from that in English-speaking countries. As stated in [19], “Catalogs play a valuable role in providing a standardized language for the activities people perform in the labor market.” While these catalogs are widely used to create and compute static values, manage labor market and educational needs, or recommend training and jobs, there is no single ground truth. According to Rodrigues (2021), one reason for this could be the fact that labor market concepts are modeled by multiple disciplines, each with a different perspective on the labor market. While there have been discussions about mapping between different standards, such as the European ESCO and the American O*NET [20], there are only limited mapping approaches between standards to date. For instance, there is no mapping between the German KldB and the Austrian AMS (we will discuss this later). This is the first gap. While there is a diverse field of different taxonomies, catalogs, and even word lists used in different institutions and for different research questions, existing tools tend to focus on only one of these perspectives, making more generic solutions difficult to implement.

For data integration, the necessity for more generic models has been discussed in the field of education, see Szabo et al. [21]. Ontologies and ontology-based methodologies have been extensively utilized. For instance, for the prediction and modeling of workshops and labor market needs, see [22], for the identification of job knowledge, see [23], but also for the analysis of particular jobs, see [24], or for the matching of educational content to generic texts, see [25]. Furthermore, these ontologies have been employed to predict the unemployment rate, as evidenced by the work of Li et al. [26]. However, these approaches have primarily concentrated on a specific labor market characteristic, such as skills, knowledge, educational content, or job classifications.

According to our best knowledge, no data warehouse or linked-data approach for labor market data has yet been proposed although some preliminary work was carried out [27], [28], [5], [29]. Thus, here we find the first gap for interdisciplinary research. In addition, we find only few works addressing the specific challenges in the analysis of labor market research data, particularly in the context of historical, qualitative, and quantitative data. Some work was carried out in the area of online social media data [30].

III. METHOD

In order to address the manifold challenges and dependencies inherent to this interdisciplinary research and data infrastructure area, it is necessary to employ a bundle of different methodological approaches. As previously discussed, despite the technical challenges that must be overcome, it is essential to address several domain-specific and research-specific questions. First and foremost, it is imperative to

house qualitative and quantitative data from a diverse range of sources, including NLP, classical social science surveys, labor market statistics, and social media data, among others. Second, we must ensure that the data is interoperable and queryable despite the lack of an overarching ontology or taxonomy for all data. Third, the data being subject to interoperability is not aggregated at the same level. For instance, we have occupations for vocational training and others linked to other domain-specific entities. Other data is linked to occupational groups, while others are linked to occupations. Fourth, quantitative data also comes with a rich assortment of metadata that describes processes and structures. These data are of great importance for the classical scientific approach in the social sciences and labor market research. Sixth, the data is often not only stored as raw data, but also in different aggregation levels, which are subject to domain-specific requirements and usually not interoperable between different data sets.

In addition to these domain-specific requirements, classical technical problems need to be solved. For example, long-term storage is needed, data protection and security needs to be addressed.

It is necessary to retrieve data from a variety of sources, including applications that retrieve data from different application programming interfaces (APIs), data analysts who work with data portals, and researchers with different privileges. It is evident that no single solution will be universally applicable; however, we will provide a comprehensive discussion of future challenges.

To tackle these these remaining interdisciplinary challenges, we will provide some methodological ideas and discuss their impact within this complex setting.

A. Structure

The data warehouse (DWH) is comprised of two primary components, see Figure 1 for an illustration. The data archive, in this instance, is the data lake, serving as the initial point of reference. The data derived from this archive is subsequently integrated into the data warehouse in a structured format.

Today, modern data stacks are mainly used with standardised connectors to operationalise ETL processes. Transformation and orchestration are usually performed by standard tools. However, due to the heterogeneity of vocational training data, custom connectors must be developed that prepare the respective data records, adapt formats, cluster and cleanse data and ensure automated data integration in such a way that they can be further used in the DWH.

The overall integration, preparation and cleansing process occurs in level 0, which is the basement of the DWH. With regard to the content of the data, it is necessary to differentiate it into three categories: raw data, L0 indicators (aggregated data that has already been processed by external systems) and a classification system. Furthermore, the documentation of the metadata is conducted in parallel. This is the point at which the expertise of domain experts is applied to the documentation.

The data is linked at the second level. As previously outlined, the GLMO classification system employs a multi-

faceted approach, encompassing a range of factors, including occupational classification (Kldb, ISCO), differentiation between gainful employment and training occupations, competence assessment (ESCO, AMS, BIBB Comp), economic indicators, and regional analysis. all taking into account the temporal course, possibly necessary anonymisation clustering (e.g. formation of occupational groups, main groups see [31]). In terms of content, the data and documentation level is of particular importance here, particularly in consideration of the findings of data analysts.

At the third level, the linked data is prepared and, if necessary, enriched, formatted for data reports or the dashboard or a data portal. In principle, this is the business intelligence (BI) level before the data is exported.

Prior to its dissemination, the data must be subjected to a final verification process and, if necessary, anonymised. At this stage, it is of particular importance to conduct plausibility and format tests. Furthermore, additional test pipelines can be developed to monitor the data analysis process.

At this level, the L2 data must be prepared by BI experts in such a way that the findings about the data discovered in Level 1 are appropriately represented, prepared and documented for the respective stakeholders. These include scientists, the CEO, and other standard users who should have access to the data, for example, in order to develop a data portal for young people for career guidance.

B. Data schemata and linked data

This section will present a selection of data schemata, which follow the from star to galaxy schema. See Figure 2 in this section for an example. Quantitative and qualitative data will be used as examples, including:

- “Datensystem Auszubildende” (DAZUBI) is a system that collates data from the vocational training statistics of the statistical offices of the Federal Government. The annual total survey encompasses data on vocational training in accordance with the Vocational Training Act (BBiG) and the Crafts Code (HwO), including trainee, contract, and examination data.
- The QuBe project is a repository of data pertaining to future qualifications and occupations. Based on economic structure models, data is forecast up to 2050. The data set contains both past data and forecast data in 726 dimensions per job, branch, region.
- Two text corpora comprising approximately nine million online job advertisements (OJAs) are available for analysis. They consist of an average of 80dim. monthly data set.
- A substantial corpus of advertisements for continuing education.
- The labor market archive.
- A diverse array of online social media data with extracted sentiments for each job
- The quali panel, which examines the structures and developments in a longitudinal perspective of company

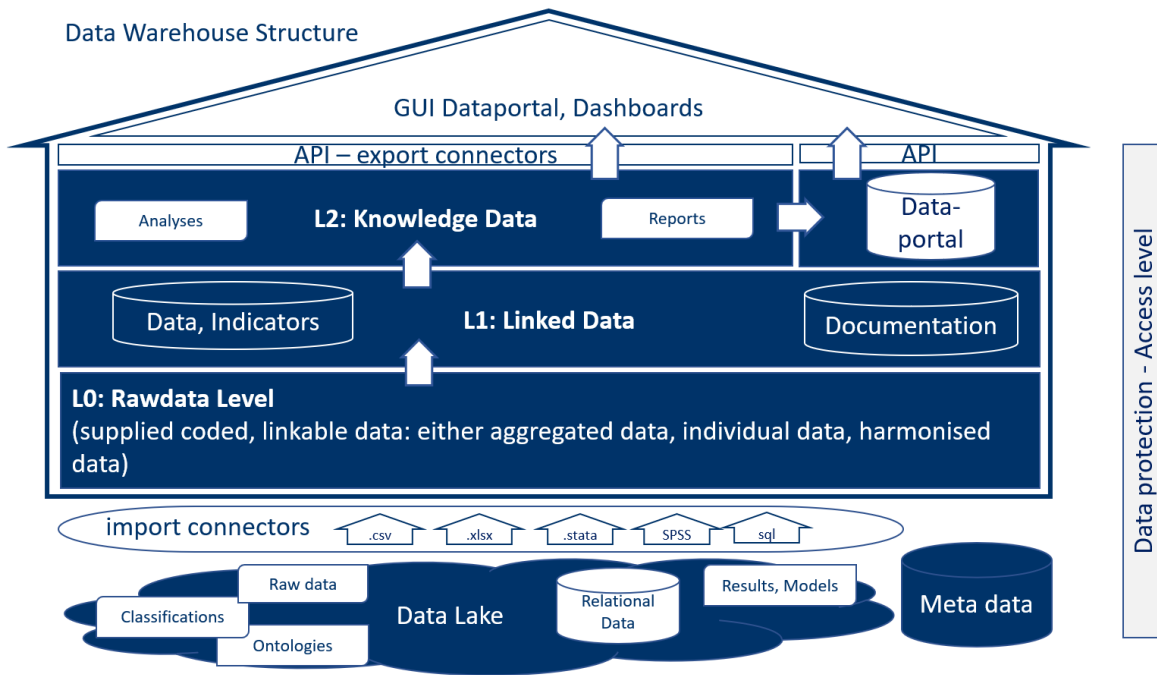


Fig. 1. Data warehouse structure: The data archive, in this instance, is the data lake, serving as the initial point of reference. The data derived from this archive is subsequently integrated into the data warehouse in a structured format.

activities and measures for training and securing skilled workers.

- The pension insurance data in the form of longitudinal and cross-sectional data contain data products for several years. The cross-sectional data products describe characteristics for reference dates (e.g. pension (e.g. pension portfolio) or reporting years (e.g. pension entries, pension histories).
- The collective training allowances for which collective agreements exist (only occupations with higher numbers of trainees)
- 42 quantitative and qualitative datasets on nursing training from 2021-2024 (only one job)
- employment surveys and 560 related indicators, like e.g. job satisfaction

It is necessary to obtain data instances that differ from one another. For instance, one instance may be used for ground truth data, while another may be employed for testing purposes. Some data will be subjected to processing. The size of the data is a significant challenge. Table I presents some information about the corresponding dimension of selected data sets.

Some data records are only collected once, whereas the majority of data records are collected and updated on an ongoing basis. In particular, for the data pertaining to job advertisements, which can be retrieved on a daily basis, an automation pipeline was designed to retrieve the new data, check for duplicates and update changes. Furthermore, an

anomaly detection system has been integrated into the data retrieval process, in order to identify any technical or content-related anomalies that may occur during the processing stage.

Due to the high dimensionality of the data and the limitations of computing resources, dimension reduction must be achieved through the preselection of the data. However, this depends on the specific use case and the objective of the data analysis. The early aggregation or elevation of data to a higher level of abstraction may result in the undesirable blurring of distinctions. Prior to the commencement of the data selection process, it is necessary to define the dimensions that are deemed to be of interest, with this selection being dependent on the specific application in question. This may entail a reduction in the number of parameters or columns, but it may also involve imposing temporal constraints or preselecting specific occupational categories, occupational sectors, or occupational groups. A preliminary selection must therefore be made by domain experts or by means of feature extraction using data science methods. At present, the selection process is still carried out manually, using prior knowledge.

In order to illustrate the methodology, we will utilize two example data sets. The first data set comprises approximately 3.5 million tweets from Twitter/X on labor market data. The second comprises the metadata of approximately 5 million YouTube videos. In Figure 1, we describe the data schema for tweets. They follow the star schema, centered around tweets. These tweets are linked data, as the `job_id` is linked to the classification of occupations (KldB) according to GLMO. Additionally, named entities are linked to CSO and GLMO.

TABLE I
SOME EXAMPLES OF DATA SETS, TYPE AND UPDATE.

Dataset	Type	Update
Dazubi	quantitative	annual
Twitter/X, YouTube, Kununu	quantitative	none (single survey)
Employment surveys (e.g. ETB)	quantitative	every 6 years
QualiPanel	qualitative	annual
QuBe	quantitative	annual
Job advertisements	quantitative	daily
Indicators derived from ETB and microcensus	quantitative	annual
Nursing training	quantitative/qualitative	annual

In Figure 2, we show the corresponding YouTube metadata schema. While this schema houses different data, it is evident that it is similarly connected to Kldb and various other ontologies, such as GLMO and CSO. This produced linked data.

The primary objective when linking the data is to identify the levels at which the data can be linked. These levels are particularly relevant in the context of the classification systems used in the GLMO of occupations at different levels. If the data records do not have an assignment to the Kldb, an ontology-based textual mapping can be employed. Alternatively, an occupation mapping via ISCO is a potential option. In some cases, datasets lack an occupation assignment but include sector information, necessitating a mapping via economic sectors. Additionally, some datasets undergo regional mapping via location parameters. In the majority of instances, individual data records are represented as individual Star Schemata. The classification system is employed to create a galaxy schema 2 through mapping.

C. Stakeholders and roles

The data warehouse stakeholders are subject to different authorisation requirements depending on whether they are accessing the individual database instances or the individual levels.

It is the responsibility of the IT architects to ensure the backup and recovery of data on the technical side. However, they are not privy to the content of the data. In contrast, data architects are responsible for the overall structuring of the data architecture, which necessitates access to the content of the data.

In addition, data engineers have access to Level 0 data, as they are responsible for developing import pipelines, harmonisation procedures and transformation matrices. In particular, data experts with a background in economics and social impact

are required for the harmonisation and alignment of marginal totals between different data sets. These domain experts then design the corresponding transformation models for the respective individual case. In subsequent work, it should be determined whether the modelling can be generalised and automated.

Those engaged in the practice of data analysis and business intelligence (BI) typically commence their training at the introductory level, which encompasses linked data and preparation for reporting, including the data portal.

IV. RESULTS

The results of our methodological approach deliver several results. First, we can demonstrate that even in this challenging interdisciplinary environment we can house linked data and make data available for querying for different stakeholders. The data can be integrated automatically via ELT pipelines with this concept and delivered for the corresponding use case or report request, as we could show for the integration of job advertisements on a daily basis. However, subsequent analysis revealed that research question 1 was overly ambitious, given the numerous challenges that were encountered. It is not possible to provide a one-size-fits-all solution. Rather, the solution must be selected on a case-by-case basis.

Second, we could identify some aspects which are subject for further research. Selecting data remained challenging, in particular because data cannot be preprocessed in endless much situations. However, ad-hoc query are not possible without further research. The necessity for analysis increases exponentially with each additional dimension of the data, indicating that manual analyses should be conducted on a select number of phenomena in the future. It is recommended that standard analysis models be developed by data scientists and that their findings be fed back into the data pool. In this context, validated results are then treated as additional multidimensional indicators analogous to weighting matrices (expert knowledge). The graphical structure of the Galaxy schema precludes the possibility of working with the storage of all abstraction levels. Consequently, it is necessary to work with abstract views and ad-hoc generations on the respective data aspect, as the overall system has to deal with limited data storage.

Third, it is necessary to consider the possibility of automatic clustering, for example at a higher occupational level due to the lack of data (sparseness of data), separately. At this time, rule-based automatic solutions are being developed.

Furthermore the consideration of qualitative data has not yet been included in the analysis and must be analysed in further work.

V. CONCLUSION AND OUTLOOK

The objective of this study was to identify the challenges and approaches to solutions that would enable the heterogeneous data landscape of vocational education and training (VET) research to be linked and reusable.

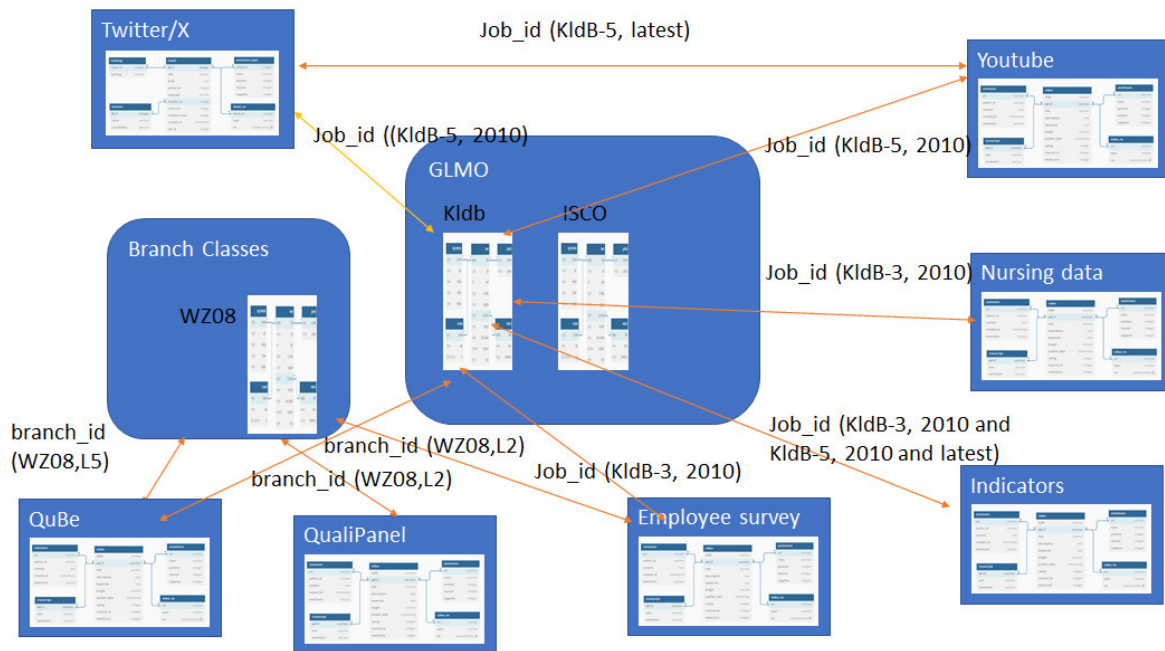


Fig. 2. Galaxy Data Schema Structure, that includes all single star and classification schemata

Although the research questions cannot be fully answered at this stage, we have identified a number of potential avenues for further investigation. These include the challenges associated with setting up a data warehouse structure for different types of data and the requirements of different stakeholders.

General challenges are how to intercept the temporal progression of data and classification systems, measurement system and dimensions change. The data can be categorised according to the level of detail provided. The following data types are to be considered: survey data, aggregated data, and linked data. A brief description of the linking process, including any pertinent notes on potential issues and expert opinions and analyses of the data.

The scientific reuse and further utilisation of all types of data requires the development of suitable procedures that ensure transparent and structured data documentation and the definition of minimum standards for the individual levels of a project dataset. We are currently developing a corresponding data management template. It would be beneficial to implement transparency and documentation, as well as data self-service, in order to facilitate the accessibility of data to the research community, without the necessity for experts in specific domains.

In light of the ever-evolving landscape of data technologies, it is prudent to adopt a modular and state-of-the-art structure wherever feasible. Migrations are a costly undertaking, requiring significant resources (time and money) and more. These include the possibility of system downtime, dissatisfaction, inconsistencies between data records, and a loss of confidence in data quality.

The harmonisation of data is still only possible with the input of experts, who are in short supply. In subsequent work, it should be determined whether the modelling can be generalised and automated.

REFERENCES

- [1] R. Dobischat, B. Käpplinger, G. Molzberger, and D. Münk, "Digitalisierung und die folgen: Hype oder revolution?" *Bildung 2.1 für Arbeit 4.0?*, pp. 9–24, 2019.
- [2] R. Helmrich, M. Tiemann, K. Troltsch, F. Lukowski, C. Neuber-Pohl, A. C. Lewalder, and B. Gunturk-Kuhl, *Digitalisierung der Arbeitslandschaften: keine Polarisierung der Arbeitswelt, aber beschleunigter Strukturwandel und Arbeitsplatzwechsel*. Wissenschaftliche Diskussionspapiere, 2016, no. 180.
- [3] J. Dörpinghaus, D. Samray, and R. Helmrich, "Challenges of automated identification of access to education and training in germany," *Information*, vol. 14, no. 10, p. 524, 2023.
- [4] J. De Smedt, M. le Vrang, and A. Papantoniou, "Esco: Towards a semantic web for the european labor market." in *Ldow@ www*, 2015.
- [5] J. Dörpinghaus, J. Binnewitt, S. Winnige, K. Hein, and K. Krüger, "Towards a german labor market ontology: Challenges and applications," *Applied Ontology*, no. 18(4), pp. 1–23, 2023.
- [6] P. Koch, W. Konen, and K. Hein, "Gesture recognition on few training data using slow feature analysis and parametric bootstrap," in *International Joint Conference on Neural Networks*, Barcelona, Spain, Jul. 2010, p. 8 pages.
- [7] A. Dutt, M. A. Ismail, and T. Herawan, "A systematic review on educational data mining," *Ieee Access*, vol. 5, pp. 15991–16005, 2017.
- [8] S. K. Mohamad and Z. Tasir, "Educational data mining: A review," *Procedia-Social and Behavioral Sciences*, vol. 97, pp. 320–324, 2013.
- [9] C. Romero and S. Ventura, "Educational data mining: A survey from 1995 to 2005," *Expert systems with applications*, vol. 33, no. 1, pp. 135–146, 2007.
- [10] J. Dörpinghaus and M. Tiemann, "Vocational education and training data in twitter: Making german twitter data interoperable," *Proceedings of the Association for Information Science and Technology*, vol. 60, no. 1, pp. 946–948, 2023.

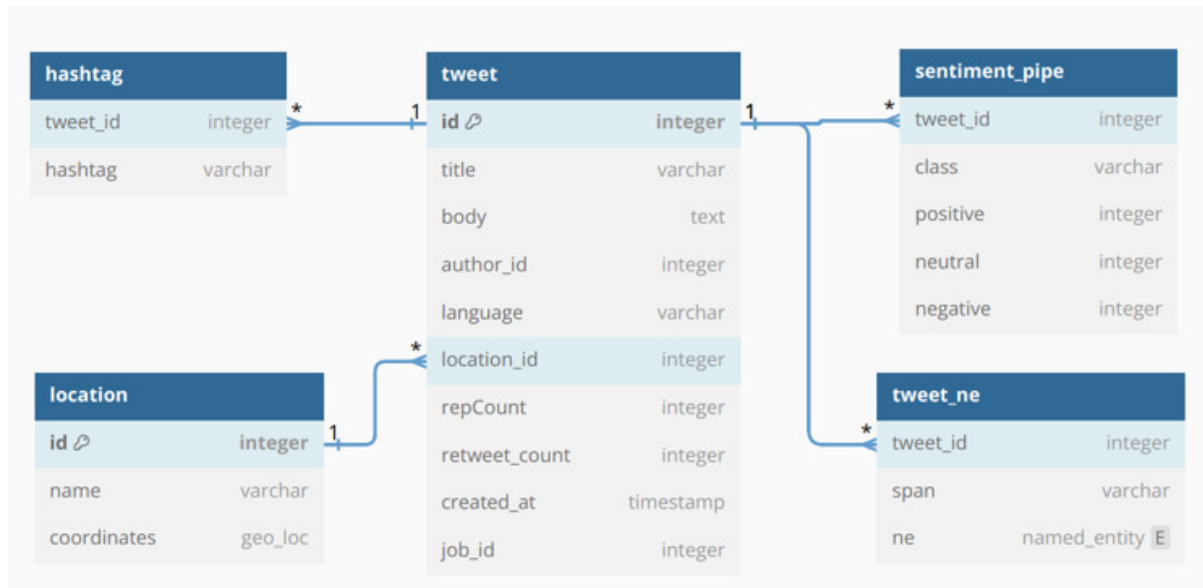


Fig. 3. Simplified star schema ERD for Twitter/X data.

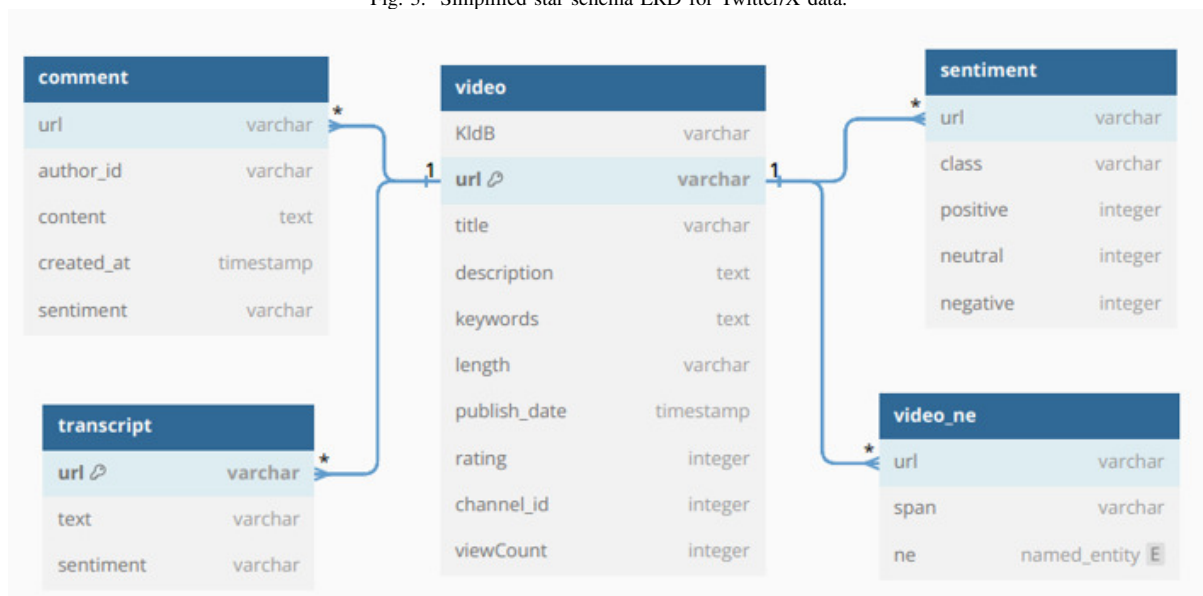


Fig. 4. Simplified star schema ERD for Youtube data.

- [11] R. M. Chang, R. J. Kauffman, and Y. Kwon, "Understanding the paradigm shift to computational social science in the presence of big data," *Decision Support Systems*, vol. 63, pp. 67–80, 2014, 1. Business Applications of Web of Things 2. Social Media Use in Decision Making. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167923613002212>
- [12] I. Moalla, A. Nabli, L. Bouzguenda, and M. Hammami, "Data warehouse design approaches from social media: review and comparison," *Social Network Analysis and Mining*, vol. 7, no. 1, p. 5, 2017.
- [13] B. Marlis, H. Buchs, and G. Ann-Sophie, "Occupational inequality in wage returns to employer demand for types of information and communications technology (ict) skills: 1991–2017," *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, vol. 72, pp. 455–482, 2020.
- [14] J. Dörpinghaus, J. Binnewitt, and K. Hein, "Lessons from continuing vocational training courses for computer science education," in *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 2*, 2023, pp. 636–636.
- [15] A. Settelmeier, F. Bremser, and A. C. Lewalder, "Migrationsbedingte mehrsprachigkeit—ein "plus" beim übergang von der schule in den beruf," *Interkulturelle und sprachliche Bildung im mehrsprachigen Übergang Schule-Beruf*, pp. 135–150, 2017.
- [16] F. Derksen and J. Dörpinghaus, "Digitalization and sustainability in german continuing education," in *INFORMATIK 2023 - Designing Futures: Zukünfte gestalten*. Bonn: Gesellschaft für Informatik e.V., 2023, pp. 1945–1953.
- [17] P. K. Ningrum, T. Pansombut, and A. Ueranantasun, "Text mining of online job advertisements to identify direct discrimination during job hunting process: A case study in indonesia," *Plos one*, vol. 15, no. 6, p. e0233746, 2020.
- [18] I. Smirnov, "Estimating educational outcomes from students' short texts on social media," *EPJ Data Science*, vol. 9, no. 1, pp. 1–11, 2020.
- [19] C. Ospino, "Occupations: Labor market classifications, taxonomies, and ontologies in the 21st century," *Inter-American Development Bank*, 2018.

- [20] S. Guru Rao, "Ontology matching using domain-specific knowledge and semantic similarity," Master's thesis, University of Twente, 2022.
- [21] I. Szabó, "The implementation of the educational ontology," in *Proceedings of the 7th European Conference on Knowledge Management, Corvinus University of Budapest, Hungary, ACL, UK*, 2006, pp. 541–547.
- [22] E. Boldyreva and V. Kholoshnia, "Ontological approach to modeling the current labor market needs for automated workshop control in higher education," in *MICSECS*, 2019.
- [23] M. Khobreh, F. Ansari, M. Fathi, R. Vas, S. T. Mol, H. A. Berkers, and K. Varga, "An ontology-based approach for the semantic representation of job knowledge," *IEEE Transactions on Emerging Topics in Computing*, vol. 4, no. 3, pp. 462–473, 2015.
- [24] M. Papoutsoglou, A. Ampatzoglou, N. Mittas, and L. Angelis, "Extracting knowledge from on-line sources for software engineering labor market: A mapping study," *IEEE Access*, vol. 7, pp. 157 595–157 613, 2019.
- [25] A. Poletaikin, S. Sinita, L. Danilova, Y. Shevtsova, and N. Dvurechenskaya, "Ontology approach for the intelligent analysis of labor market and educational content matching," in *2021 International Symposium on Knowledge, Ontology, and Theory (KNOTH)*. IEEE, 2021, pp. 50–55.
- [26] Z. Li, W. Xu, L. Zhang, and R. Y. Lau, "An ontology-based web mining method for unemployment rate prediction," *Decision Support Systems*, vol. 66, pp. 114–122, 2014.
- [27] T.-P. Liang and Y.-H. Liu, "Research landscape of business intelligence and big data analytics: A bibliometrics study," *Expert Systems with Applications*, vol. 111, pp. 2–10, 2018.
- [28] T. Avdeenko and M. Bakaev, "Modeling information space for decision-making in the interaction of higher education system with regional labor market," in *2014 12th International Conference on Actual Problems of Electronics Instrument Engineering (APEIE)*. IEEE, 2014, pp. 617–623.
- [29] A. Fischer and J. Dörpinghaus, "Web mining of online resources for german labor market research and education: Finding the ground truth?" *Knowledge*, vol. 4, no. 1, pp. 51–67, 2024.
- [30] B. Batrinca and P. C. Treleaven, "Social media analytics: a survey of techniques, tools and platforms," *Ai & Society*, vol. 30, pp. 89–116, 2015.
- [31] M. Tiemann, H.-J. Schade, R. Helmrich, A. Hall, U. Braun, and P. Bott, "Berufsfeld-definitionen des bibb auf basis der klassifikation der berufe 1992," *Schriftenreihe des Bundesinstituts für Berufsbildung Bonn*, vol. 105, no. 1, p. 57, 2008.

An autonomous vehicle in a connected environment: case study of cyber-resilience

Guillaume Hutzler*, Hanna Klaudel*, Witold Klaudel^{†‡}, Franck Pommereau* and Artur Rataj[‡]

* IBISC, Univ. Evry, Université Paris-Saclay, France

Email: {guillaume.hutzler, hanna.klaudel, franck.pommereau}@univ-evry.fr

[†] SafeTech Cybernetics, Palaiseau, France, Email: witold.klaudel@outlook.fr

[‡] IRT SystemX, Palaiseau, France, Email: artur.rataj@irt-systemx.fr

Abstract—As the advancing autonomy of vehicles requires increasing assistance from the surrounding infrastructure, it becomes clear that the potential for cyberattacks necessitates a sophisticated implementation of resilience, capable of detecting and responding to both internal and external threats. Therefore, threat analysis and risk assessment, including careful modelling of resilience, are essential to prepare against cybersecurity risks. In this context, we extend our method of an automatic discovery of cost-ranked cyberattack scenarios by monitoring/fallback mechanisms. We then demonstrate that this extension allows an analysis of a realistic resilient model of cybersecurity aspects of a level 2 autonomous vehicle in a connected environment.

Index Terms—security evaluation, formal modelling

I. INTRODUCTION

AUTONOMOUS vehicles, still in the experimental stage, are far, from a technical and legal point of view, from established complex resilient systems such as e.g. energy networks. However, recent progress demonstrates that an autonomous vehicle will be part of large infrastructure systems, including elements such as manufacturer’s diagnostic supervision, dealer authorisation, various services such as geospatial or even road infrastructure. Obviously, this shapes resilience goals. We reflect this by immersing, in our case study, an autonomous vehicle into its environment and modelling attacks that often cross the boundary between these two parties.

Technological advances have expanded the cyberattack surface of distributed information systems, and specifically critical ones such as autonomous vehicles, as they have become more complex and also more connected with the external world. This resulted in sophisticated implementations of resilience understood as active defence, capable to detect attacks and react to them [1]. Careful modelling of resilience mechanism may be necessary, given that a number of new applications of complex distributed information systems concern critical systems [2]. This is the case of the autonomous vehicle studied in this paper.

In order to meet these new expectations, we extended our framework, called SCORE [3], which is devoted to build a suitable automata model of attack propagation in the system and automatically discover complex cyberattack scenarios using abstract cost criteria. The new extension adds to SCORE a complementary monitoring/fallback mechanism implementing resilience on an abstraction level compatible with a Model-Based Systems Engineering (MBSE) architectural diagram of

hardware and software. We then show that together with an extraction of security traits from a heterogeneous non-security oriented MBSE data, this allows analysing a realistic resilient model of the cybersecurity aspect of an autonomous vehicle. The automated security analysis of architectural traits provided by SCORE greatly optimises an analyst’s work by allowing her/him to focus less on engineering aspects and more on purely cybersecurity related concepts like abstract costs of unitary compromise of software components.

Identifying threats to distributed information systems is challenging as they consist of many different entities exposed to a wide range of cyberattacks. A cyberattack is understood here as a sequence of unitary actions taken by an *attacker* to take control over some components of the system. This sequence starts with one or more entry points and ends with the loss of integrity of some system components leading in turn to a damage targeted by the attacker. Identifying cyberattacks is crucial for optimising both an architecture and cybersecurity features to meet an acceptable level of risk. It is necessary to calculate likelihood of cyberattacks and thus contribute to the estimation of the overall risk level.

When it comes to interpreting the concept of likelihood, many national cybersecurity agencies (e.g., NIST [4], ANSSI [5]) issue loosely defined recommendations and defer final decisions for its rating scale and methods to experts. The choice of scale and methods for assessing the strength of cyber protections is still an open research problem, e.g. [6]. In SCORE we lean towards the ANSSI approach translating the likelihood into an inversely proportional cost. We compute the cost of the entire attack as a sum of costs of all its unitary attacks, assuming that experts can provide the scale and value of costs associated with breaking the cyber protections of the systems under analysis.

Paper structure: We first provide a related work and a comparison between SCORE method and the existing ones (Sec. II). Then, in Sec. III we remind the essential characteristics of SCORE and Sec. IV extends it with resilience features such as monitoring and fallback. In Section V we apply SCORE to the cyber resilience risk assessment of a realistic electronic onboard architecture of autonomous vehicle. We develop a propagation model of this case study and discuss the compliance of the chosen security protections wrt the acceptable risk level.

TABLE I
FEATURE AND MODEL SIZE COMPARISON.

Method	Automatic calculation				Active defence	Case study		
	Network access	Coalition	Peer position	MITM position		Edges	Nodes	Hard. nodes
Static AT or ADT [7], [8]	n.a. ^a	+	n.a. ^a	n.a. ^a	-	<20	<20 ^b	n.a. ^c
Architecture to AG [9], [10]	+	+	-	-	+	<20	<20 ^d	n.a. ^c
Architecture to AT and AG [11]	+	+	-	-	+	<20	<20 ^d	n.a. ^c
Architecture/VDB to AG [12]	+	+	-	-	-	<100	<20 ^d	<20
Archit./Assert./VDB to AG [13]	-	-	+	-	-	<1000	<1000 ^b	n.a. ^c
Archit./VDB to AG/BDD [14]	-	-	-	-	-	>50k	>50k ^d	n.a. ^c
Our approach SCORE	+	+	+	+	+	<10k	<100 ^d	<100

^a Because of lack of architecture.

^b Attacker's sub-goals [8].

^c Architecture is not layered.

^d Software nodes or *de-facto* software (a hardware node with no internal software structure given).

II. CONTRIBUTION AND RELATED WORK

The first non-trivial approaches to vulnerability models were static Attack Trees (AT) [7], [15], [16], where nodes represented logical operations like an AND gate showing the necessity for several parallel breaches before an attack can proceed — we call it a coalition. While static, these trees could be created by various manual procedures [17] e.g. motivated by a security property of interest. Later, ATs were extended to include defence mechanisms to form Attack-Defence Trees (ADT) [8].

Due to the increasing complexity of distributed systems' architecture, the overhead of identifying attack propagation paths between architectural components grew exponentially. This posed a burden for security analysts and presented a potential source of human error. The necessity of devising an automated analysis of the topology of a complex distributed system became obvious. [9] proposed a method of creating a coincidence matrix between a number of architectural components like servers or routers, producing thus a general Attack Graph (AG) to be processed e.g. by a model checker in order to find possible attack sequences. We see that method as one of the first attempts at reusing a traditional MBSE architectural data for cybersecurity, then extended in numerous ways [18], [19], [20].

In SCORE, starting from hardware and software system architectures annotated by cybersecurity protection features, we build an attack propagation model in terms of a network of automata, each automaton modelling a software module of the system. Each automaton evolves according to a set of attack propagation rules computing the cost for each state change as a function of states of contributing neighbouring software components. The automaton state represents the type of software component compromise, called its status, and may be:

- nominal, not yet compromised, thus unable to propagate;
- active malware, which has largest spectrum of possible propagation,
- passive (called bad data), which may only send a corrupted data through the network, or
- non-available, which represents a component completely disabled by the attacker, so that even the attacker cannot

use it anymore.

We compare our framework to representatives of different families of methods with respect to the attack structures and the size of systems under analysis in Tab. I.

Here we present an overview of features allowed by SCORE:

- all software components are divided into three classes: user, root and kernel, corresponding in our approach to system processes;
- we employ *functional interactions* representing producer-consumer relations between user software components;
- we consider *system interactions* allowing to model basic operating system relations, like a cheap attack from a kernel to a process it manages or an attack against a kernel via a network interface controlled by that kernel;
- thanks to the automatic calculation of possible flow of interactions through the hardware layer, attacks may take advantage from the so called attacker's position, which can be Man-In-The-Middle (MITM) or Peer, depending on the relationship between the attacking component with the interaction it attacks;
- we also classify certain elements in order to decrease the number of free security parameters like unitary attack costs; this eases the work of a security analyst and increases the manageability of a model; a unitary attack cost consists of a protocol and a component breach;
- thanks to the modelling of the system as a network of automata, we are able to model synchronous attacks (coalition) in which an attacker possibly propagates in a non-sequential way.

See that a component status and an attacker's position, combined with the class of a software component, can form together a rich Cartesian product whose tuples (like *malware* performing a *MITM* attack against a *user component*) can be used in security properties, like for example arrays of unitary costs. We see the abstract nature of such tuples, as opposed to concrete attack descriptions (like a worm X uses a vulnerability Y against a component version Z). This high level of abstraction results from the motivation behind the method: estimation of general resilience based on architectural traits and not an identification of concrete vulnerabilities/exploits as understood e.g. within the CVE

database [21]. This is a substantial difference to attack graph-based vulnerability/exploit search tools [22], [10], [11] often connected to popular vulnerability databases [23], [12], or analysers of the source code of concrete IoT devices [24].

Functional and system interactions are a unique trait of our approach. They are similar to the notion of trust in [13]. Each interaction may be routed through the network or through the kernel if located on the same hardware component. In our approach, the routing possibilities of each interaction are pre-calculated statically on the basis of all possible network transmission paths between interaction partners. These paths may be seen as a Boolean function of the interaction accessibility for the attack propagation; it means that if it evaluates to false the attack cannot propagate through the interaction.

As [9] already noticed, solving attack graphs can be much more numerically intensive than solving attack trees. Thus various methods combining trees and graphs [25], [26] or assuming the criterion of monotonicity [27] *i.e.*, that a unitary attack may not decrease further attacking capabilities. Very large networks have been analysed thanks to the latter property [14] combined with binary decision diagrams [28]. However, using the above mentioned abstract classification with a very limited number of classes and a simple abstract model behaviour of a software component, a modern model checker on a fast hardware was able to solve in a reasonable time an involved MBSE model (Sec. V) without the monotonicity assumption. We take advantage of the latter which allows more realistic attack scenarios. For example, an attacker can disable a router or trigger a monitoring system, both potentially limiting the possibility of further propagation.

III. FORMAL DESCRIPTION

We shortly remind the basic elements of SCORE before introducing its extension with monitoring and fallback. More detailed description of the SCORE approach may be found in [3].

The SCORE system specification starts with the definition of the hardware and software architectures and the respective mapping. Each software component exposes interfaces, which are necessary for its normal functioning but can be abused by an attacker to penetrate the system. Interface exposure depends on functional complexity of the system but also on hardware architecture constraints such as resource sharing and network routing. It is assumed that an attacker can pass through any interface but possibly with different abstract costs. The definition of these costs is another part of the system specification that is used to generate the propagation model.

Fig. 2 and 3 show an example of the hardware and software layers. Hardware components are connected by undirected network links while software components are connected by directed functional interactions encapsulated in communication protocols, *e.g.* HTTP for web applications. Software interaction endpoints are attached to software components through software ports called *roles*, *e.g.* client and server for HTTP protocol. Software components are hosted by hardware ones and functional interactions are routed through network links.

SCORE assumes that each software component is assigned to exactly one hardware component. A privileged software component (kernel) may also manage its hardware host, including subordinate software components within the same hardware host and network or intra-system routing rules. The hardware link types correspond to the communication link types such as for example, Ethernet or CAN network; they connect components through hardware ports. Hardware components are managed by operating systems (kernels). Routing rules must ensure a physical realisation of all specified functional interactions. Often, due to technical and organisational constraints, routing rules are more permissive than necessary, which can result in the creation of additional attack opportunities.

The propagation model generation goes through an intermediate construction, called the *visibility graph*, which results from the synthesis of the architecture and focuses only on the information flow allowing propagation of the attacker in the system. The visibility graph is a directed graph indicating which software component (node) can propagate its corrupted status to its successors. Each interaction between two nodes is directed and gives rise to an edge in the visibility graph connecting a node to the role of its target node. Edges are labelled by the *attacker position* relative to this interaction:

- *peer* if the attacking node is a functional peer within that interaction or if the edge points to a system role;
- *mitm* (man in the middle) if the attacking node is on a routing path taken by the interaction;
- *side* if the routing rules merely allow the attacking node to see the target role but the node is not in *peer* or in *mitm* positions.

In order for the interaction to be effective, at least one communication routing path should be available, *i.e.*, the statuses of router nodes on this path should be different from non-available. This is expressed statically by a Boolean formula obtained from the architecture when generating the visibility graph.

To obtain an analysable propagation model P with secrets, the visibility graph is completed by some auxiliary information such as thresholds, secrets and roles' categories and the abstract costs of unitary attacks, set by experts. P is essentially composed of a network of automata $\{A_1, \dots, A_{|N|}\}$ (one for each node $n \in N$ of the visibility graph), each of them being in its current state in $L = \{\mathcal{F}, \mathcal{N}, \mathcal{B}, \mathcal{M}\}$ meaning respectively functional, not available, bad-data or malware, and secrets in S refer to security keys protecting communication sessions, which can be stolen from the nodes where they are stored. All the automata in $\{A_1, \dots, A_{|N|}\}$ have identical structure, *i.e.*, the same states and transitions, see Fig. 1, but different transition firing conditions. The last may be complex and depend on the type of the software component and on its connectivity with other components in the visibility graph. Detailed definitions of transition conditions are provided in [3].

A. Propagation model dynamics

A *configuration* of the propagation model P with secrets is the state of the network of automata and the state of secrets,




Fig. 1. General shape of the automaton A_n of node $n \in N$.

which is subject to evolve during the execution. More formally, a configuration c is a pair (\vec{q}, \vec{s}) , where $\vec{q} \in L^{|N|}$ is the vector of states of the automata and $\vec{s} \in \mathbb{B}^{|S|}$ the current Boolean value of secrets. For example, $([\mathcal{F}, \mathcal{F}, \mathcal{F}, \mathcal{M}], [0, 0, 1])$ is a configuration of a model with four nodes (software components) and three secrets; first three nodes are functional (\mathcal{F}) and the last one is corrupted (\mathcal{M}), the third secret key is stolen (globally known).

The firing of a transition from a configuration to another corresponds to a transition in some automaton A_n in the network; it computes the abstract cost for the attacker to reach the target configuration. A transition resulting in a status change of node n in automaton A_n has a unitary cost and depends on the status of n itself and on the status of its inputs (predecessors) $p \in \text{pre}(n)$ in the visibility graph. The impact in terms of cost of a predecessor p on n may vary depending on the role r of n to which p is connected; it depends on the category and type of r , and on the attacker position labelling the edge from p to r in the visibility graph. Some transitions in A_n may need a set of predecessors (a coalition) to have particular statuses in order to be enabled. The theft of secrets impacts the cost calculation of the transition by eliminating the costs of breaking the communication protocols' security of the interactions that they protect. Stolen secrets are visible globally; this corresponds to an omniscient attacker that has a global knowledge of the system (or at least the parts that they control). The overall cost of a transition in A_n is the sum of the cost of breaking communication protocols and breaking node's role protections, or equals to the cost of stealing secrets by the node itself when its status is already \mathcal{M} (malware).

Given a configuration of the propagation model, several transitions may potentially reach a successor configuration; however, we assume in SCORE that only the transition of minimal cost between a pair of configurations is considered in the definition of the system dynamics. This choice greatly improves performance while being consistent with SCORE's objective of discovering attacks with maximum likelihood, *i.e.*, minimum cost. Formally, this gives the following definition of the semantics:

Definition 3.1: The semantics of P with secrets S is a transition system $(\text{Config}_P, \rightarrow, c_0)$ where Config_P is the set of all configurations reachable from an initial configuration $c_0 = (\vec{q}_0, \vec{s}_0)$ by executing transitions defined as follows:

A transition from (\vec{q}, \vec{s}) to (\vec{q}', \vec{s}') with cost κ , denoted by $(\vec{q}, \vec{s}) \xrightarrow{\kappa} (\vec{q}', \vec{s}')$, exists if there exists an enabled transition $t_n = (\eta, \eta')$ in some automaton A_n with minimal cost κ , *i.e.*, $\vec{q}[n] = \eta$, $\vec{q}'[n] = \eta'$, and for all $i \in [1..|N|]$, $i \neq n$, $\vec{q}'[i] = \vec{q}[i]$, which updates accordingly the secrets, *i.e.*, $\vec{s}' = \text{update}_n(t_n, \vec{s})$. \diamond

Attack discovery in the SCORE propagation model identifies sequences of interface crossings that lead from the initial configuration, with one or more attackers positioned in software components, to an undesirable target configuration. Among the huge number of possible attacks, SCORE selects, using model checking queries, that having a minimal cost or a cost under some fixed bound, representing the maximum likelihood needed to maintain the acceptable level of risk.

IV. EXTENSION WITH RESILIENCE

Resilience is the ability of a system to operate under adverse conditions or stress, even if in a degraded mode, while maintaining essential operational capabilities, and to recover to a nominal operational mode. In the initial version of SCORE [3] only the mechanisms of access control, isolation and redundancy were proposed, which is not sufficient to cover most of the resilience requirements. In particular we were not able to model degraded mode nor recovery. As in distributed information systems the recovery process is complex and often includes human decision, in this paper we decided to focus on rising degraded modes. This was also needed by the application to the autonomous vehicle we had in mind.

In order to take into account a part of resilience requirements we introduce a *monitoring* concept, which considers for each target configuration $c \in \text{Config}_P$ a possible fallback one. More formally, we define a function $\text{fallback}: \text{Config}_P \rightarrow \text{Config}_P$ indicating for each configuration a corresponding fallback one, and an associated function $\mu: \text{Config}_P \rightarrow \mathbb{N}$ representing the cost of bypassing the monitoring in the target configuration.

Intuitively, the semantics of a propagation model P with resilience policy $P' = (P, \text{fallback}, \mu)$ is then obtained by replacing each transition $c_1 \xrightarrow{\kappa} c_2$ existing in the semantics of P , by two transitions: one with cost κ leading from c_1 to a fallback configuration $\text{fallback}(c_2)$ and another with cost $\kappa + \mu(c_2)$ from c_1 to the initial target configuration c_2 . This means that from c_1 with cost κ we may only reach the fallback configuration of c_2 , while c_2 remains reachable up to an additional cost $\mu(c_2)$.

Definition 4.1: The semantics of a propagation model with resilience policy $P' = (P, \text{fallback}, \mu)$ is a transition system $(\text{Config}_{P'}, \leftrightarrow, c_0)$, where $\text{Config}_{P'}$ is the set of all configurations reachable from c_0 by executing transitions in \leftrightarrow , defined as follows: if $c_1 \xrightarrow{\kappa} c_2$ with some cost κ , then we have $c_1 \xrightarrow{\kappa + \mu(c_2)} c_2$ and $c_1 \xrightarrow{\kappa} \text{fallback}(c_2)$. \diamond

Note that the set of all configurations of P' and P are identical, $\text{Config}_{P'} = \text{Config}_P$. However, the accessibility of certain configurations may be modified.

V. CASE STUDY

In this section, we introduce our case study representing the on-board electronics of a Level 2 autonomous connected vehicle, which means that the vehicle can control its speed and direction in some specific situations but the human driver must be able to regain full control of the vehicle at any time.

A. Presentation

As represented in Fig. 2 and 3, the use case architecture is composed of two main parts: the **Vehicle** part comprising

- the three vehicle control domains:
 - **Power Train**, which covers engine and gearbox control (if applicable);
 - **Body**, which covers vehicle access control, cabin lighting, headlights, windscreen wipers and air conditioning;
 - **Chassis**, which includes brakes, steering, ultrasound and cluster.
- **ADAS** (Advanced Driver Assistance Systems), which includes front camera, lidar, front and rear radars, and assist mode switch;
- **Communication**;
- **Multimedia**, and
- **Central Gateway**, separating the critical parts of the vehicle from the Internet connectivity and from the multimedia, and allowing the navigation interacting with the external world. The central gateway switch **gateSwitch** filters the network connections and the central gateway unit **gateUnit** contains software components in charge of central diagnostics and software updates and the proxying activities between critical and exposed parts of the vehicle,

and a simplified representation of the **External Infrastructure** comprising

- **Internet** with a content provider and cellular network,
- **Dealership** with the capability of vehicle diagnostics and vehicle software update,
- **OEM** (car maker) with the central management of vehicle software and navigation map delivery,
- **GNSS** (Global Navigation Satellite System) responsible of the vehicle geographical positioning.

Both, the **Vehicle** and the **External Infrastructure**, have their hardware and software architectures, as shown in Fig. 2 and 3, respectively.

Concerning the hardware architecture of the vehicle, each of the three vehicle domains: **Power Train**, **Body** and **Chassis** have its own domain controller, which communicates with the components inside the domain using a separate CAN network. The **ADAS** domain is built around an Ethernet switch **adasSwitch** allowing the exchange of a large volume of data between the **ADAS** controller **adasCtrl** and all the domain components. All these four domains mentioned above communicate through inter-domain CAN network **interd**. Three complementary Ethernet links between **ADAS**, **Power Train** and **Chassis** domains allow to exchange of a large data

necessary for advanced functionalities (**ADAS**); for example, displaying the image from the rear camera **rCamera** on the cluster **cluster**. The central gateway area **Central Gateway** consists of two components: a switch **gateSwitch** in charge of network traffic filtering and redirection, and a gate unit **gateUnit** in charge of central management functions and proxying activities necessary for supplementary information flow verification and separation. The switch **gateSwitch** is connected by Ethernet to **gateUnit**, **adasSwitch**, multimedia/navigation **mMedia**, the communication unit **commUnit**, and to the dealer diagnostic devices **dealer** during the dealer intervention. The communication unit **commUnit** connects the vehicle to the Internet through the cellular network. The gate unit **gateUnit** has also a supplementary link to the inter-domain CAN network **interd**. The body controller **bodyCtrl** communicates via radio link with the vehicle's access card **card**, and the multimedia and navigation unit **mMedia** communicates via Bluetooth with a smartphone **phone**, for example to stream music to the vehicle. Of course, the smartphone is connected to the Internet.

The above architecture may seem suboptimal, but it reflects the current real-world situation of automakers who prefer to reuse existing solutions to extend the functionality of their products. These incremental transformations require a very detailed analysis of the cyber risks generated by the newly introduced interactions between initially independent subsystems.

Concerning the software architecture, we distinguish the following categories of components:

- kernels, shown with a dashed border,
- components with root privileges, shown in light pink with a brown border,
- user space components, shown in light blue with a dark border.

Each software component has a name starting with a capital letter and its hosting hardware component indicated on a dark background in the lower part. The interactions are indicated using links and sometimes using pairs of labels having the same colour, especially when it would be too complicated to trace lines, but there is no semantic difference between these two representations. A hardware component may host at most one kernel. Kernels have no depicted explicit interactions but implicitly, each kernel manages system interactions with all software components hosted by the same hardware component. Kernel-less hardware components can only host software components having root privileges.

B. Attacker entry points

Potential attacker entry points to the system are also indicated in the architecture definition. They appear in orange in Fig. 2 with a Trojan icon. We consider the following entry points for attackers:

- **intHacker** is located on the Internet;
- **dealHacker** is located in the dealership;
- **oemHacker** is located in the intranet of the carmaker (OEM);

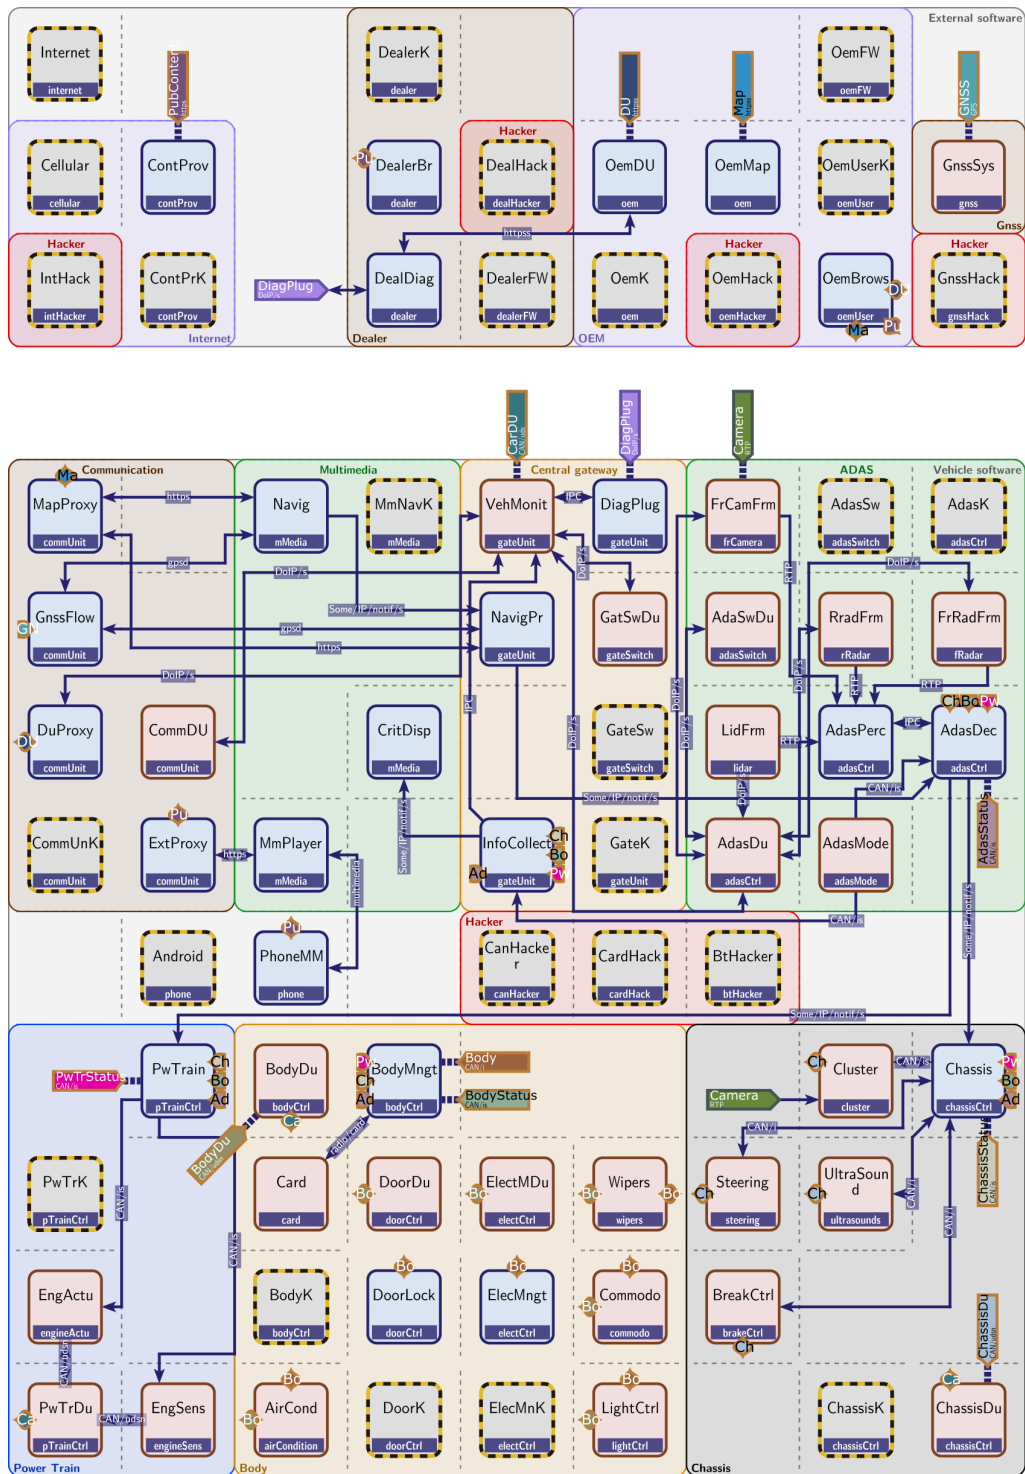


Fig. 3. Software architecture.

consists in defining at least one active hacker, *i.e.*, such that its status is different from functional \mathcal{F} ; usually we set it to malware \mathcal{M} . The target configuration consists in indicating which components have to be compromised, *i.e.*, the status of one or more of them should be different from \mathcal{F} .

D. Analysis

Our study is a part of the risk analysis process such as defined in standards like ISO 21434 [30] or EBIOS Risk Manager [31]. It starts from the definition of so called *feared events* associated here with vehicle functionalities and reveals a harmful breach of them. Each feared event is given a severity degree and an accepted level of risk. Both together allow to estimate for each feared event the maximal accepted likelihood of attacks, which can trigger this event.

In this paper, we assume that the above stage is already provided and the list of feared events is known along with the related likelihood, represented by its inverse, *i.e.*, the acceptable cost, as shown in Tab. II. We also assume that the unitary costs are already defined for all protocols and all the types of software components used in our architectures, taken from libraries provided by experts. Intuitively, the costs may be divided in three categories: weak (under 16), intermediate (between 15 and 25) and strong (between 26 and 50). Our contribution consists then in looking for potentially dangerous attacks capable to trigger the mentioned feared events, computing their costs in order to either confirm that the level of risk remains acceptable, or to propose modifications leading to better protections. The analysis we present in the following is developed in order to illustrate the method. It is of course partial and provides only a small subset of attacks, which should be included in a complete industrial cybersecurity risk assessment.

TABLE II
FEARED EVENTS AND ACCEPTABLE COSTS VS COMPANY RISK
MANAGEMENT POLICY

	feared event	acceptable cost
1	vehicle theft	80
2	ADAS sabotage	80
3	driver disturbance	40

Our analysis will proceed as follows: For each feared event *feared*, we will look for possibly several attacks depending on the initial and target configurations. More precisely, we choose a set of meaningful attacker entry points and a set of target software components to be compromised in order to trigger *feared*. This leads to several cases represented in Tab. III and IV together with their calculated minimal costs and the corresponding shortest path.

In order to compute these attack paths of minimal costs we set for each case the initial configuration in the propagation model expressed in UPPAAL and use the model checker. It is worth to mention that the UPPAAL formula allowing to find the minimal cost of an attack of a given target may be prohibitively long to compute due to the size of the model

and the combinatorial explosion of the number of reachable configurations. Actually, using the "brute force" is not efficient (we stopped the computation after a few hours) and we need to use hints to make this computation feasible. The method consists in first assuming that such an attack exists and use the model checker to confirm that it is under a given estimated cost. The difference is that such a restricted request is generally much faster. It's even faster as this estimated cost gets closer to the searched minimum cost. Usually, we start from largest estimations and refine them by dichotomy until obtaining an acceptable computation time, usually under 30m. Then, still under the constraint of the estimated cost, we search the attack of the minimum cost, which take usually a few minutes.

For example, in the case 1a in Tab. III we set first the status of *IntHacker* to \mathcal{M} and that of all the other components to \mathcal{F} in the initial configuration. The UPPAAL formula

$$\text{inf}\{\text{DoorLock.B}\}: \text{costs} \quad (1)$$

looks for the minimal value of variable *costs* of an attack path reaching the target configuration where the component *DoorLock* has status \mathcal{B} . However, as mentioned above, it is not efficient without constraining the cost under some estimated bound. To find such a bound, we check under different cost constraints the UPPAAL CTL formula

$$\text{E}\langle\langle \text{DoorLock.B} \rangle\rangle \quad (2)$$

saying that there is a path reaching eventually the configuration where the component *DoorLock* has status \mathcal{B} . If the cost constraint is close to the minimum, the formula (1) becomes efficient and it allows to find the attack of the minimal cost quite fast (in few minutes, under 10m). The minimal calculated cost is 100 and the obtained path is shown in Tab. III together with the cost of each step.

We consider six cases for the vehicle theft depicted in Tab. III. In the first three, the attacker comes from the Internet, and looks for compromising the functionality of the access to the vehicle and its start, which are managed by the software component *BodyMngt*. In all these cases the attack path passes through the update functionality. The second case shows that if *BodyMngt* is compromised, the opening of the doors becomes almost costless. The third case shows a malware installation on *BodyMngt*, which results in a persistent access to the vehicle, but at the price of a higher cost. In the fourth case the objective is as in the previous ones but when the vehicle is connected to the dealer diagnostic tool via *DiagPlug* so that the attack passes through the dealership. In the fifth case, the attacker enters through Bluetooth interface. This attack is costly because in our setting Bluetooth does not allow any vehicle opening functionalities. In the sixth case, the attacker is initially located in the dealership, attacks first the diagnostic device, then the body management in order to install a malware to be used when the vehicle is outside the dealership.

Concerning ADAS sabotage, see Tab. IV, we consider two cases with the same attacker entry point from the Internet. In the first case the target is the break control *BreakCtrl* and the second the ADAS decision function *AdasDec*.

TABLE III
ATTACKS, COSTS AND PATHS FOR FEARED EVENT 1.

	entry/status	target/status	cost	path (comp/status:cost)
1a	IntHacker/M	DoorLock/B	100	IntHacker/M → DuProxy/B:50 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:20 → DoorLock/B:5
1b	IntHacker/M	BodyMngt/B	95	IntHacker/M → DuProxy/B:50 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:20
1c	IntHacker/M	BodyMngt/M	105	IntHacker/M → DuProxy/B:50 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/M:30
1d	IntHacker/M	BodyMngt/M	105	IntHacker/M → DealDiag/B:55 → DiagPlug/B:5 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/M:30
1e	BtHacker/M	BodyMngt/B	120	BtHacker/M → MmPlayer/M:25 → Navig/B:10 → NavigPr/B:20 → AdasDec/B:15 → BodyMngt/M:50
1f	DealHack/M	BodyMngt/M	90	DealHack/M → DealDiag/B:30 → DiagPlug/B:5 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:30

TABLE IV
ATTACKS, COSTS AND PATHS FOR FEARED EVENTS 2 AND 3.

	entry/status	target/status	cost	path (comp/status:cost)
2a	IntHacker/M	BreakCtrl/B	80	IntHacker/M → DuProxy/B:50 → VehMonit/B:20 → ChassisDu/B:5 → BreakCtrl/B:5
2b	IntHacker/M	AdasDec/B	85	IntHacker/M → MapProxy/B:50 → NavigPr/B:20 → AdasDec/B:15
3a	IntHacker/M	MmPlayer/B	45	IntHacker/M → PhoneMM/B:40 → MmPlayer/B:5
3b	IntHacker/M	MmPlayer/B	45	IntHacker/M → ExtProxy/M:40 → MmPlayer/B:5

For the possibility of driver disturbance, see also Tab. IV, we consider two attacks starting from the Internet and targeting the multimedia player MmPlayer, for example to force the maximal speaker volume. The first attacks through the phone connected to the vehicle, while the second through the HTTP proxy; both have the same cost.

The calculated minimal costs of all the above attacks are

acceptable from the risk assessment point of view, however one may observe a high dependability on the resistance of the proxy functionalities. As a consequence, a further reinforcing could be useful. A possible solution could be provided by a monitoring with a reaction such as a fallback. This will be illustrated in the next section.

E. Monitoring and fallback

In this section we provide the system with monitoring and fallback of three proxies: ExtProxy, DuProxy and MapProxy. Each of them may of course be lured but at the price of a supplementary cost.

We assume that an attempt to attack ExtProxy may lead to switch it off (*i.e.*, force its status to non-available \mathcal{N}) meaning that the vehicle loses the access to the Internet multimedia content, we call this fallback action Fb_1 . The remaining two proxies are more critical, so we assume that the fallback action Fb_2 cuts off the whole communication between the vehicle and the external world, *i.e.*, the whole Central gateway part is forced to switch off. In the propagation model this is represented by forcing GateK and GateSw to status \mathcal{N} .

TABLE V
ATTACKS, COSTS AND PATHS WITH MONITORING AND FALLBACK FOR FEARED EVENT 1.

	entry/status	target/status	cost	path (comp/status:cost)
1a	IntHacker/M	DoorLock/B	130	IntHacker/M → DuProxy/B:80 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:20 → DoorLock/B:5
1b	IntHacker/M	BodyMngt/B	125	IntHacker/M → DuProxy/B:80 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:20
1c	IntHacker/M	BodyMngt/M	135	IntHacker/M → DuProxy/B:80 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/M:30
1d	IntHacker/M	BodyMngt/M	105	IntHacker/M → DealDiag/B:55 → DiagPlug/B:5 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/M:30
1e	BtHacker/M	BodyMngt/B	120	BtHacker/M → MmPlayer/M:25 → Navig/B:10 → NavigPr/B:20 → AdasDec/B:15 → BodyMngt/M:50
1f	DealHack/M	BodyMngt/M	90	DealHack/M → DealDiag/B:30 → DiagPlug/B:5 → VehMonit/B:20 → BodyDu/B:5 → BodyMngt/B:30

In Tab. V and VI we show the same combinations of entry points and targets for the three mentioned above feared events. As expected, the costs of attacks which are able to avoid

proxies are unchanged. All the attack paths passing through the proxies are more expensive. However, new paths having lower costs appear, like those of cases 2a' and 2b', which pass through the dealership.

In order to find the attacks luring the monitoring and calculate their costs, we need to assume that either the vehicle is not plugged on the dealer diagnostic tool (cases 1a, 1b, 1c, 1e, 2a and 2b), or that it is not connected to the phone (case 3b). It may be obtained using the UPPAAL formula, for example $\text{inf}\{\text{MmPlayer.B and Android.F and PhoneMM.F}\}$: costs for the case 3b.

TABLE VI
ATTACKS, COSTS AND PATHS WITH MONITORING AND FALLBACK FOR
FEARED EVENTS 2 AND 3.

	entry/status	target/status	cost	path (comp/status:cost)
2a	IntHacker/M	BreakCtrl/B	110	IntHacker/M → DuProxy/B:80 → VehMonit/B:20 → ChassisDu/B:5 → BreakCtrl/B:5
2a'	IntHacker/M	BreakCtrl/M	105	IntHacker/M → DealDiag/B:55 → DiagPlug/B:5 → VehMonit/B:20 → ChassisDu/B:5 → BreakCtrl/M:15
2b	IntHacker/M	AdasDec/B	115	IntHacker/M → MapProxy/B:80 → NavigPr/B:20 → AdasDec/B:15
2b'	IntHacker/M	AdasDec/B	105	IntHacker/M → DealDiag/B:55 → DiagPlug/B:5 → VehMonit/B:20 → AdasDu/B:5 → AdasDec/B:15
3a	IntHacker/M	MmPlayer/B	45	IntHacker/M → PhoneMM/B:40 → MmPlayer/B:5
3b	IntHacker/M	MmPlayer/B	95	IntHacker/M → ExtProxy/M:70 → MmPlayer/B:5

As shown above, adding monitoring and fallback improves system protection. However, because this introduces a new cyber attack surface, triggering fallback mode may constitute a new target for the attacker. This is visible in our case study as an attack from the Internet taking as the objective to rise fallback Fb_2 is possible for a rather low cost. The UPPAAL formula

$\text{inf}\{(\text{GateK.N and GateSw.N})\}$: costs

returns the minimum cost of 50.

F. Redundancy

In this section we illustrate a situation where the attacker is able to pass through the routing restrictions on the switches of Central Gateway and ADAS, *i.e.*, GateSw and AdasSw, for example if it succeeded to get the network access key. This allows it to attack communication protocols between perception (AdasPerc) and various sensors such as lidar, radar

or camera, in order to force perception to create a fake scene, *i.e.*, a false image of the road situation. Usually, these protocols are weakly protected.

As these sensors are redundant in the sense that they provide partly the same information, in our case the attacker has to compromise at least two of them, which makes it more costly.

Concerning the network access key, we consider that the attacker steals it from the OEM Diagnostic and Update Center (OemDU).

We are able to detect such an attack by placing the attacker in the Internet, *i.e.*, in IntHacker and by looking for paths compromising ADAS perception with the following formula:

$\text{inf}\{(\text{AdasPerc.B and DiagPlug.F})\}$: costs

The condition DiagPlug.F is added to confirm that the vehicle is not connected to the dealer diagnostic tool. The corresponding path and cost is given in Tab. VII.

TABLE VII
ATTACK, COST AND PATH FOR SENSOR COMPROMISE.

entry/status	target/status	cost	path (comp/status:cost)
IntHacker/M	AdasPerc/B	110	IntHacker/M → OemDU/key-theft:35 → ExtProxy/M:55 → AdasPerc/B:20

The first step in the path consists in stealing remotely the protection key without changing the status of the OEM Diagnostic and Update Center. Then, the attacker installs malware on the External Proxy luring its monitoring but taking profit from the stolen access key. Finally, the attacker compromises synchronously two interactions between the Perception and two of its sensors to produce fake data. This is costly but less expensive than the installation of a malware on the Perception.

VI. CONCLUSION AND PERSPECTIVES

We show how this method can be used for cybersecurity risk assessment of large critical systems. It can also suggest directions for optimising both passive architectural and active cyber protections in order to achieve an acceptable level of risk. We expect that such mechanisms will eventually become compulsory for resilient critical systems, and in particular for autonomous vehicles.

We notice, that certain threat scenarios depend on the functionality of software components which is unrelated to security and is thus not covered by our security-focused approach. For example, in the case of a vehicle, it may be interesting to know whether the vehicle is stationary or running, whether the navigation system is connected to the Internet or not, or whether only certain ADAS functionalities, such as the autopilot, is active or not. These extensions are of course possible, but may degrade computational efficiency. The challenge for our future work will be to find an acceptable trade-off between these aspects.

We see two improvements related to unitary attack costs: their automatic synchronisation with the CVE database [32]

and an enhanced method of their cumulation into the total scenario expense. The former would improve on resilience thanks to timely dissemination of vulnerability information. The latter would take into account elements such as partial observability of the system from the perspectives of different actors.

ACKNOWLEDGMENT

This work was supported by the French government as part of the “France 2030” program, within the framework of the SystemX Technological Research Institute.

REFERENCES

- [1] A. Clark and S. Zonouz, “Cyber-physical resilience: Definition and assessment metric,” *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 1671–1684, 2017. doi: 10.1109/TSG.2017.2776279
- [2] N. Leveson, N. Dulac, D. Zipkin, J. Cutcher-Gershenfeld, J. Carroll, and B. Barrett, “Engineering resilience into safety-critical systems,” in *Resilience engineering*. CRC Press, 2017. doi: 10.1201/9781315605685-12 pp. 95–123.
- [3] G. Hutzler, H. Klauudel, W. Klauudel, F. Pommereau, and A. Rataj, “Automatic discovery of cyberattacks,” in *IEEE CSR*, 2024, to appear.
- [4] S. Quinn, N. Ivy, M. Barrett, L. Feldman, G. Witte, and R. Gardner, “Identifying and estimating cybersecurity risk for enterprise risk management,” 2021. doi: 10.6028/NIST.IR.8286A <https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8286A.pdf>.
- [5] “Digital risk management,” French Cybersecurity Agency, 2024, <https://cyber.gouv.fr/en/digital-risk-management>.
- [6] S. Gupta Bhol, J. Mohanty, and P. Kumar Pattnaik, “Taxonomy of cyber security metrics to measure strength of cyber security,” *Materials Today: Proceedings*, vol. 80, pp. 2274–2279, 2023. doi: 10.1016/j.matpr.2021.06.228 SI:5 NANO 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214785321046009>
- [7] S. Mauw and M. Oostdijk, “Foundations of attack trees,” in *Information Security and Cryptology-ICISC 2005*. Springer, 2006. doi: 10.1007/11734 pp. 186–198.
- [8] J. Arias, C. E. Budde, W. Penczek, L. Petrucci, T. Sidoruk, and M. Stoelinga, “Hackers vs. security: attack-defence trees as asynchronous multi-agent systems,” in *International Conference on Formal Engineering Methods*. Springer, 2020. doi: 10.1007/978-3-030-63406-3_1 pp. 3–19.
- [9] R. Ritchey and P. Ammann, “Using model checking to analyze network vulnerabilities,” in *IEEE Symposium on Security and Privacy*, 2000. doi: 10.1109/SECPRI.2000.848453 pp. 156–165.
- [10] S. Jajodia, S. Noel, and B. O’berry, “Topological analysis of network attack vulnerability,” *Managing Cyber Threats: Issues, Approaches, and Challenges*, pp. 247–266, 2005. doi: 10.1145/1229285.1229288
- [11] M. Ge, J. B. Hong, W. Guttman, and D. S. Kim, “A framework for automating security analysis of the internet of things,” *Journal of Network and Computer Applications*, vol. 83, pp. 12–27, 2017. doi: 10.1016/j.jnca.2017.01.033
- [12] C. Hankin, P. Malacaria *et al.*, “Attack dynamics: an automatic attack graph generation framework based on system topology, capec, cwe, and cve databases,” *Computers & Security*, vol. 123, p. 102938, 2022. doi: 10.1016/j.cose.2022.102938
- [13] O. Sheyner and J. Wing, “Tools for generating and analyzing attack graphs,” in *International symposium on formal methods for components and objects*. Springer, 2003. doi: 10.1007/978-3-540-30101-1_17 pp. 344–371.
- [14] K. Piwowarski, K. Ingols, and R. Lippmann, “Practical attack graph generation for network defense,” in *Computer Security Applications Conference*. IEEE Computer Society, 2006. doi: 10.1109/ACSAC.2006.39. ISSN 1063-9527 pp. 121–130. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/ACSAC.2006.39>
- [15] B. Schneier, “Attack trees,” *Dr. Dobbs’s journal*, vol. 24, no. 12, pp. 21–29, 1999.
- [16] B. Kordy, S. Mauw, S. Radomirović, and P. Schweitzer, “Attack–defence trees,” *Journal of Logic and Computation*, vol. 24, no. 1, pp. 55–87, 06 2012. doi: 10.1093/logcom/exs029. [Online]. Available: <https://doi.org/10.1093/logcom/exs029>
- [17] D. M. Kienzle and W. A. Wulf, “A practical approach to security assessment,” in *Proceedings of the 1997 workshop on New security paradigms*, 1998. doi: 10.1145/283699.283731, pp. 5–16.
- [18] M. S. Barik, A. Sengupta, and C. Mazumdar, “Attack graph generation and analysis techniques,” *Defence Science Journal*, vol. 66, no. 6, p. 559, 2016. doi: 10.14429/dsj.66.10795
- [19] H. S. Lallie, K. Debattista, and J. Bal, “A review of attack graph and attack tree visual syntax in cyber security,” *Computer Science Review*, vol. 35, p. 100219, 2020. doi: 10.1016/j.cosrev.2019.100219. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1574013719300772>
- [20] K. Kaynar, “A taxonomy for attack graph generation and usage in network security,” *Journal of Information Security and Applications*, vol. 29, pp. 27–56, 2016. doi: 10.1016/j.jisa.2016.02.001. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214212616300011>
- [21] MITRE, “Common weakness enumeration,” 2023, <https://cwe.mitre.org/data/index.html>.
- [22] O. Sheyner, J. Haines, S. Jha, R. Lippmann, and J. Wing, “Automated generation and analysis of attack graphs,” in *IEEE Symposium on Security and Privacy*, 2002. doi: 10.1109/SECPRI.2002.1004377, pp. 273–284.
- [23] I. Chokshi, N. Ghosh, and S. K. Ghosh, “Efficient generation of exploit dependency graph by customized attack modeling technique,” in *Advanced Computing and Communications*. IEEE Computer Society, 2012. doi: 10.1109/ADCOM.2012.6563582, pp. 39–45. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/ADCOM.2012.6563582>
- [24] Z. B. Celik, P. McDaniel, and G. Tan, “Soteria: Automated {IoT} safety and security analysis,” in *USENIX Annual Technical Conference*, 2018. doi: 10.48550/arXiv.1805.08876, pp. 147–158.
- [25] J. Hong and D.-S. Kim, “Harms: Hierarchical attack representation models for network security analysis,” 2012. doi: 10.4225/75/57b559a3cd8da
- [26] J. B. Hong and D. S. Kim, “Towards scalable security analysis using multi-layered security models,” *Journal of Network and Computer Applications*, vol. 75, pp. 156–168, 2016. doi: 10.1016/j.jnca.2016.08.024.
- [27] P. Ammann, D. Wijesekera, and S. Kaushik, “Scalable, graph-based network vulnerability analysis,” in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, 2002. doi: 10.1145/586110.586140, pp. 217–224.
- [28] R. E. Bryant, “Graph-based algorithms for boolean function manipulation,” *Computers, IEEE Transactions on*, vol. 100, no. 8, pp. 677–691, 1986. doi: 10.1109/TC.1986.1676819
- [29] G. Behrmann, A. David, and K. G. Larsen, “A tutorial on UPPAAL,” in *LNCS*, vol. 3185. Springer, 2004. doi: 10.1007/978-3-540-30080-9_7, pp. 200–236.
- [30] “Road vehicles, Cybersecurity engineering,” International Organization for Standardization, Geneva, CH, Standard, 2021.
- [31] “Ebios risk manager,” French Cybersecurity Agency, 2024, https://www.ssi.gouv.fr/uploads/2019/11/anssi-guide-ebios_risk_manager-en-v1.0.pdf.
- [32] “Common vulnerabilities and exposures,” MITRE, 2024. [Online]. Available: <http://cve.mitre.org>

Learning from the COVID-19 Pandemic to Improve Critical Infrastructure Resilience using Temporal Fusion Transformers

Jakob Jenko 0009-0007-8655-4838 XLAB Ljubljana, Slovenia	Joao Pita Costa* 0000-0001-5745-1302 XLAB Ljubljana, Slovenia joao.pitacosta@xlab.si	Daniel Vladušič 0000-0001-6032-8575 XLAB Ljubljana, Slovenia	Urban Bavčar 0009-0009-4046-0312 ELES Ljubljana, Slovenia	Radoš Šabarkapa 0009-0007-5202-4572 EKC Beograd, Serbia
---	--	---	--	--

Abstract—During the COVID-19 pandemic, traditional demand prediction models drastically failed mostly due to altered consumption patterns. Accurate forecasts are essential for ensuring grid stability.

This paper analyzes the performance of the Temporal Fusion Transformer (TFT) model during the COVID-19 pandemic aiming to build resilient demand prediction models. Through detailed analysis, we identify which features may contribute to improved performance during large-scale events such as pandemics. During lockdowns, consumption patterns change significantly, leading to substantial errors in existing demand prediction models.

We explore the impact of features such as mobility and special day considerations (e.g., lockdown days) on enhancing model performance. We demonstrate that periodic updates on a monthly basis make the model more resilient to changes in consumption patterns during future pandemics.

Moreover, we show how improvements in prediction accuracy translate to real-world benefits, such as enhanced grid stability and economic advantages, including reduced energy waste. Additionally, we discuss the implications for energy-critical infrastructure, considering disruptive scenarios like future pandemics.

I. INTRODUCTION

As the integration of renewable energy sources into power grids intensifies, and the accuracy of energy demand predictions becomes increasingly crucial. Effective energy management requires that energy production aligns closely with demand to minimize the losses often associated with over-production. Therefore, reliable and precise demand forecasts are essential for optimally adjusting production levels.

Moreover, extreme situations like global pandemics can drastically change consumption patterns overnight, underscoring the importance of having adaptable and responsive demand prediction models. These models must quickly incorporate new data and adjust to shifting consumption dynamics to ensure energy efficiency and grid stability. By doing so, they help maintain a balance between production and demand, preventing inefficiencies and promoting sustainable energy use. Moreover, the discussion on the trustworthiness and certification of AI systems and, in particular, of neural networks as in [1] is essential to move forward with a fast changing technological landscape.

Maintaining the alignment between production and demand is critical for grid stability, as deviations can lead to significant issues, including a drop in the grid frequency below 50 Hz, potentially causing grid collapse or separation. In scenarios where demand exceeds supply, gas-powered peaker plants are typically utilized to provide the necessary additional capacity. However, some of this demand can also be mitigated through the use of pumped-storage hydroelectricity, which contributes stored, often renewable, energy back into the grid. In situations where it is not feasible to meet high demand exclusively through increased production, load shedding [2] is implemented as a controlled process to prevent total grid failure. This involves selectively disconnecting parts of the grid—such as entire neighbourhoods—to reduce the overall electrical load, ensuring that the grid does not exceed its capacity.

Accurate forecasting models are indispensable to grid management, particularly in anticipating and responding to demand surges. This capability becomes even more crucial during unforeseen critical events, such as pandemics, which can abruptly and drastically alter usage patterns. Effective models must rapidly adapt to new consumption patterns, providing timely forecasts that reflect current consumption trends to maintain grid stability. The pertinence of machine learning methods in the study of energy efficiency in particular pandemic scenarios gains much from the data collected in the most recent COVID-19 global pandemic. The intersection of energy efficiency and artificial intelligence (AI) has gained unprecedented significance, as the crisis reshaped global energy consumption patterns and highlighted the urgency of sustainable practices. The most recent machine learning methods emerged as key enablers in adapting to these changes.

This paper discusses how AI-driven solutions can be instrumental in optimizing energy use during the pandemic, ensuring efficient operations while addressing the environmental challenges exacerbated by the health crisis. By examining AI's role in mitigating energy consumption in a time of fluctuating demand and promoting sustainable practices in the face of adversity, this analysis illuminates the critical role of

technological innovation in navigating the energy challenges posed by COVID-19 taking into consideration mobility. We have a closer look at the mobility data in the context of a pandemic, based on the data collected during the COVID-19 incidence between 2020 and 2022. Particularly, we look at the number of unique connections to the cell tower in Slovenia measuring how much people migrate. The consumption curve can describe the behaviour of people in regard to mobility. Taking into account that Telecom data is usually expensive, we compare the relevance of that data for the forecasting model in relation to the usage of labels for special days (e.g., considering lockdown days as holidays).

The research question addressed in this paper regards if the input of mobility data is comparable to the input of special days, particularly during a pandemic scenario learning from the experience (and data) from the COVID-19 pandemic. In particular,

- Can we improve model accuracy and performance during an impactful large-scale event (such as a pandemic) with additional features (e.g. mobility and other special day features)?
- Can we make the model more resilient by periodically updating on a monthly basis
- How do the improvements in prediction accuracy reflect in the real world? Benefits to grid stability, benefits to economic aspects as less energy is wasted etc.
- What should energy Critical Infrastructure (CI) take into account in disruptive scenarios like future pandemics?

The main contribution of this paper is a new methodology based on the Temporal Fusion Transformer (TFT), and its initial evaluation, which shows how past energy consumption, weather forecast and energy-saving features can impact the prediction of energy consumption.

Results will be presented in Section III-A, where we demonstrate the model's performance in predicting Serbian national consumption. This analysis will validate and illustrate how the initial version of our model performs against the established EKC model. Following this comparison, we will build upon this model in subsequent subsections, focusing on data from Slovenia.

The COVID-19 pandemic has significantly altered consumption patterns due to increased home stays, underscoring the need to integrate mobility data into forecasting models for enhanced accuracy. This approach is supported by [3], demonstrating the effectiveness of incorporating mobility features from publicly available Google data into their predictive models, significantly improving forecast precision. While authors in [3] were able to demonstrate this phenomenon across multiple states and continents using mobility data for the US and EU, our study is limited to Slovenia due to data constraints. However, we anticipate that this phenomenon will be applicable to other cultures and states, as observed in the cited paper. In line with these findings, our methodology involves deriving a "mobility factor" from data provided by a Slovenian national telecommunication provider. By analyzing the number of unique connections to cell towers, we can infer

mobility patterns: fewer cell connections typically indicate that residents are staying home while connecting to multiple cells suggests movement to different locations. The total number of unique connections across all cells in a given area reflects the overall mobility, serving as a valuable predictor in our models.

We enhance model performance during periods of rapid consumption changes (e.g., lockdowns) by incorporating a "special day flag." This flag is activated on weekends, holidays, or days with enforced curfews/lockdowns. The advantage of this flag lies in its simplicity and availability for day-ahead forecasting scenarios, providing a straightforward yet effective method to account for unusual consumption patterns. This paper discusses the implemented state-of-the-art deep learning models, with particular focus on the TFT approach [4]. Moreover, we build on decision tree-based models such as XGBoost [5] and CatBoost [6]. Linear regression was used as a baseline, and the energy data was sourced from ELES and EKC, represented in this paper by the respective coauthors. The models are further refined by including mobility data as an additional input, which is expected to bolster their accuracy, particularly during periods like the pandemic when typical consumption patterns are disrupted. By adapting these advanced models to incorporate new, relevant data inputs, we propose a new standard in forecasting precision, ensuring optimal energy management even in the face of significant behavioural shifts induced by global crises.

This research work builds on [7] and [8], in the context of CIs as addressed by the Horizon Europe project SUNRISE building resilience in cases of unforeseen events, such as pandemics. It particularly focuses on the needs of CIs (like railway, water distribution operators, hospitals, etc.), however, it is clearly evident that the main dependency of all CIs is electricity. Ensuring stable electricity availability (stability of the network) is dependent on many factors, the most prominent being patterns of energy consumption and in recent years, renewable energy production (typically solar production).

II. METHODOLOGY

A. Temporal Fusion Transformer and self attention

Transformer-based models have been shown to surpass traditional architectures like recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) in performance, making them an attractive option for a wide range of applications. The methodology we consider in this study reapplies the TFT approach [9], leveraging the transformers' architecture and self-attention mechanisms inherent to this architecture. The TFT model accommodates the input of various variables using a variable selection network (VSN), which assesses the significance of each input. This system enhances the influence of more impactful inputs while diminishing the effects of less relevant and noisy data. Based on these evaluations, inputs are combined and subsequently processed further. The merged inputs are sent in a LSTM, used to make sense of temporal relations between the time stamps recurring to past and future covariates. The subsequent phase applies the static enrichment layer, which is particularly beneficial

when dealing with numerous categories, such as price, carbon emissions, and load demand. This layer enhances the model’s handling of such diverse classes. Following this, the temporal self-attention mechanism comes into play, where the model prioritizes (i.e. focuses attention on) the most critical time positions. This mechanism is crucial for identifying both long-term and short-term dependencies within the observed and known time-varying inputs. In the training phase, the model minimizes the loss function by tweaking the weights, which in the VSN amplifies the impact of significant features while suppressing the noisy ones. Similarly, the attention mechanism concentrates on time positions that have a substantial effect on predictions. A clear example is in energy data analysis, where the model gives precedence to past weekends to enhance the accuracy of weekend consumption forecasts.

To train the models we’ve used Optuna, a hyper parameter optimization framework [10], and for the loss function we utilised Quantile loss [11].

TABLE I
TABLE OF HYPER-PARAMETERS FOR THE 24H MODEL

Hyper-parameter	Value
attention head size	32
dropout	0.28
hidden size	92
hidden continuous size	64
learning_rate	0.001
batch_size	64
lstm_layers	2
max_encoder_length	24

The optimal hyper-parameter set for our 24h model is shown in Table I. Note that for the larger 168h model, more parameters are required. Please refer to the original TFT paper [9] for those values. Given these hyper-parameters total number of trainable parameters was roughly 900,000.

B. Data Collection and Processing

The study utilizes historical weather and energy data that have been publicly shared by Transmission System Operators (TSOs) in two countries, Slovenia and Serbia, in the context of the SUNRISE project. This data collection underpins the research, providing a foundational dataset for the TFT-based analysis in this paper. The research employs historical weather measurements sourced from the open Meteostat platform [12] and solar irradiation data from Open Meteo [13]. These sources provide freely available data for research purposes, with the exception of historical forecasted weather data, which was procured in bulk from OpenWeatherMap [14] for the cities of Belgrade and Ljubljana. This forecasted data is crucial for accurate evaluations as weather predictions are updated several times a day, and using historical forecasts helps to prevent data leakage.

Specifically from the Open Meteo database, only shortwave radiation data was utilized as it was not available from other sources. This dataset is accessible via the Open Meteo API and is licensed under the open-source Creative Commons 4.0

license. It includes comprehensive meteorological data for Slovenia, detailing parameters such as temperature, relative humidity, dew point per square meter, apparent temperature, precipitation levels, rainfall, snowfall, snow depth, atmospheric pressure, surface pressure, cloud coverage, wind speed, wind direction, wind gusts, and notably, shortwave radiation.

Conversely, the Meteostat dataset comprises measured meteorological data for Serbia, which includes temperature, dew point temperature, actual humidity, precipitation, snowfall, wind direction, wind speed, peak wind gusts, pressure, daily sunshine duration, and weather condition codes. This dataset is pivotal for the Serbian energy consumption forecast tool and the benchmarking forecasting model. It is also publicly accessible and can be retrieved via API.

The electricity consumption dataset encompasses historical data on national electricity usage. This data was sourced from the ENTSO-E Transparency Platform [15] and enhanced with data from EKC. They provided baseline modelled forecasts for energy based on demand/consumption in megawatts (MW) with an hourly resolution. This comprehensive data collection allows for highly accurate comparisons, facilitated by using the same training cutoff date. Training and cutoff dates are specified in each experiment separately. If not, the training start date was the start of 2017 for "long" models and 2019 for the rest. The 2019 cutoff is related to the mobility data cutoff date.

Additionally, mobility data was supplied by Telekom Slovenije, the national telecommunications provider in Slovenia. This dataset tracks the daily number of unique connections to each cell within the network, excluding connections from hosted users. Each cell tower is divided into multiple cells, and the number of unique connections per cell serves as an indicator of mobility. Essentially, the more frequently users move and change cell towers, the higher the total count of unique connections, which in turn provides a measure of the mobility factor.

Figure 1 illustrates the variations in the mobility factor from 2019 to 2023, with annotations for the three lockdown periods. The data indicates that the reductions in mobility during the summer holidays are similar to those observed during the first lockdown. Notably, the initial lockdown had the most profound impact on mobility, with each subsequent lockdown having a progressively lesser effect; the third lockdown, in particular, shows a minimal influence on mobility patterns.

To properly analyse the performance of the models during the times of altered consumption patterns, we plan to focus on the year 2020. All the data used was sampled at 1 sample per hour and was normalised using z-score normalisation.

C. Implementation

To develop the service and train the model, we utilized Python 3.10, supplemented by several key libraries aimed at data manipulation and mathematical operations. Specifically, we used Numpy [16] for numerical computations, Pandas [17] for data analysis, Matplotlib [18] for plotting graphs, and Scipy [19] for additional scientific computations. For the deep




Fig. 1. Mobility between 2019 and 2023

learning components, we employed PyTorch [20], a popular framework for deep learning applications. Alongside, we used PyTorch Forecasting [21], an extension of PyTorch designed specifically for time series forecasting. This combination of tools provided a robust environment for developing complex predictive models efficiently.

D. Data Challenges

Efforts have been concentrated on understanding the basic processes involved in each of the use cases, especially in the context of the COVID-19 pandemic, which presented unprecedented challenges to urban management worldwide. The pandemic severely disrupted daily routines and behaviors as individuals who were exposed to or contracted the virus had to isolate themselves, inhibiting their ability to perform normal activities. Communities enforced social distancing measures to mitigate transmission risks. These widespread disruptions contributed to significant societal and economic impacts, including a substantial loss of life. The adjustments made during the pandemic have highlighted the importance of adaptive strategies in managing public health crises.

The dataset preliminarily consist of aggregated and fully anonymized data concerning people's activity levels, as recorded by the telecommunication provider. This primarily includes the number of individuals present in a specified area (e.g., a municipality) at a given time. The data is aggregated both spatially (to the level of municipalities) and temporally (to hourly intervals), ensuring that it is impossible to extract any privacy-sensitive information. This approach is similar to the methodologies used in the Google COVID-19 Community Mobility Reports and Apple COVID-19 Mobility Trends, which provided public access to mobility data during the pandemic. However, these sources are no longer updated and suffered from limited regional coverage and resolution. The

current dataset aims to fill these gaps by providing more detailed and continuously updated information.

E. Evaluation

To assess the performance of our models, we utilized both the Mean Absolute Error (MAE) and the Mean Absolute Percentage Error (MAPE) as metrics. MAPE is particularly valuable for its intuitive interpretation, making it easier to understand forecasting accuracy. Our primary focus was on predicting load demand, a standard benchmark that enables comparison with other methodologies. Additionally, we extended our analysis to include predictions on energy prices and carbon emissions, demonstrating the versatility and broad applicability of our models in various contexts. This comprehensive evaluation helps highlight the models' effectiveness across different domains.

$$MAE = \left(\frac{1}{n}\right) \sum_{i=1}^n |y_i - x_i| \quad (1)$$

The Mean Absolute Error (MAE) is beneficial for quantifying the actual prediction errors, which can be particularly useful when analyzing individual signals. However, for comparing performance across different models or tasks, the Mean Absolute Percentage Error (MAPE) tends to be more suitable. This metric, expressed as a percentage, provides an intuitive measure of a model's accuracy.

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - x_i}{x_i} \right| \quad (2)$$

III. RESULTS

To demonstrate that the TFT model can improve upon current state-of-the-art methods, we will compare it to EKC and ELES in-house models currently used in production settings. Next, we will analyze the performance of the TFT model during the first COVID-19 lockdown in Slovenia and investigate the effects of additional features on the TFT, as well as on gradient boosting methods like XGBoost and CatBoost. Finally, we will perform a detailed analysis of the feature importance of the TFT model.

A. Comparing TFT and EKC models

First, we will demonstrate how our base model compares against the EKC, to predict Serbian national consumption. The evaluation was performed between 1.1.2022 and 15.1.2023, and the model was trained on data between 1.1.2019 and 31.12.2022. In Table III we present the comparison between the XLAB TFT model and the EKC model. As we can observe in Table II, we have used future forecasted weather as well as a bigger input window size of 168h instead of 24h. Here we demonstrate that by increasing the amount of information we managed to gradually improve the model's performance.

In Table III we can further explore the results above in Table II, but in higher detail. Both the MAPE and MAE metrics improved by up to 20%, accompanied by a reduction in standard deviation and a lower maximum error.

TABLE II
PERFORMANCE COMPARISON FOR 2022 BETWEEN EKC BASELINE AND XLAB TFT MODEL

Model	MAPE [%]	Forecasted Weather	Window Size [h]
EKC_baseline	2.34	No	24
XLAB_TFT 1	2.18 (-7 %)	No	168
XLAB_TFT 2	1.93 (-17 %)	Yes	24
XLAB_TFT 3	1.87 (-20 %)	Yes	168

TABLE III
MAPE AND MAE COMPARISON BETWEEN EKC BASELINE AND FOR BEST PERFORMING TFT VARIANT

Metric Model	MAE		MAPE [%]	
	EKC	XLAB	EKC	XLAB
mean	92.24	73.58 (-20 %)	2.34	1.87 (-20 %)
std	91.00	67.53	2.075	1.65
max	3540.00	860.00	29.60	20.05

By extending the input window size to 168 hours, we achieved improved results, however, this also led to an increase in training time. In the next section, we will utilise the 24h model, as it offers a better ratio between time to converge and performance.

B. Comparing Iterative TFT and ELES Models

In the case of ELES, the specific training cut-off date was not known, making direct comparisons potentially unfair since one model might have more recent knowledge updates than the other. To address this issue, we developed an iterative model that updates its knowledge on a monthly basis, ensuring a fairer comparison. The model was trained on data from 1.1.2018 up to 31.12.2019. The last 14 days of this period were used for validation during the training loop. Subsequent months were used for testing, progressively incorporating more data into the training set.

In regards to the iterative model, Table IV shows the comparison with the ELES in-house model and the improvements we achieve with the approach proposed in this paper. As we can observe in the Table for 2020, the ELES model outperformed our approach, whereas, for 2021 and 2022, we were able to improve their approach by up to 10.59 % on average. The reason for the decreased performance in 2020 could be attributed to multiple factors where one of them could be related to the model not having enough data and of course big change in consumption patterns due to the COVID-19 pandemic.

Overall, an iterative approach inherently results in a more resilient model. When new consumption patterns emerge, they will be automatically incorporated into the next model iteration. An example is expressed in the data of the COVID-19 pandemic. Since no lockdowns or similar events were present in the training data, we cannot expect our model to make accurate predictions during the first COVID-19 lockdown. Even if we do not pass any information about mobility or lockdown, the model should adapt to new consumption

TABLE IV
MAE AND MAPE METRICS FOR ITERATIVE LEARNING (2020-2022)

Year	Metric model	MAE		MAPE [%]	
		ELES	XLAB	ELES	XLAB
2020	mean	41.77	43.26 (+3.57%)	2.81	2.92 (+3.91%)
	std	37.94	40.37	2.56	2.72
	max	689.00	730.00	30.66	29.55
2021	mean	43.46	37.34 (-13.87%)	2.73	2.37 (-13.1%)
	std	42.82	34.26	2.58	2.10
	max	319.00	300.00	21.57	20.00
2022	mean	43.62	40.07 (-8.14%)	2.92	2.69 (-7.88%)
	std	44.32	35.75	3.21	2.49
	max	1214.00	340.00	91.97	47.22

patterns in the next iteration. Alternatively, we could assume that the lockdown day consumption pattern is similar to that of a weekend or a holiday. Setting a flag that would treat lockdown days as holidays would not confuse the model as much as it would without such a flag. In the next Section III-C, we will include this feature to inspect its impact.

C. Model Performance During COVID-19 Lockdown

In this section, we will utilise the very same models as in Section III-B and dive deeper into analysis during COVID-19 lockdowns in 2020 for Slovenia. By adding new features, we demonstrate how we can improve the prediction capabilities of existing models.

The Table V is calculated based on Table VI. It demonstrates that the relative difference between model predictions during lockdown and non-lockdown periods. The first two models, *ELES* and *XLAB base long* are the same models as in Table IV. The remaining models are similar but include additional features; for example, *XLAB specday long* incorporates a special day feature. The models *base*, *mobility*, and *mobility specday* follow a similar pattern but use less training data, specifically starting from January 1, 2019—approximately one year less than the other models. Consequently, this may result in poorer performance for these models, as they were trained with roughly 14 months of data by the first lockdown.

In Figure V we can observe a noticeable decrease in performance during the first COVID-19 lockdown. After removing the outliers, the change is roughly **32.66%**, which is in line with the literature in [3] observing a similar impact. While the analysis of the change is not relevant, it demonstrates a pattern that models utilising either mobility or special days have lower differences i.e. performing better during the first lockdown.

Most relevant is the analysis of Table VI, where we are focusing on the first column *lockdown 1* presenting performance during the first lockdown period in Slovenia. The interval can be visually observed in Figure 1.

The best-performing model is *XLAB spec day long*. Based on the comparison with its baseline, it yielded a relative improvement of roughly 18.10% and 10.13% compared to the model that was trained on data from 2019 onward. The improvement compared to the ELES model was less than 1%.

TABLE V
RELATIVE DIFFERENCES BETWEEN MAPE DURING LOCKDOWN
NO-LOCKDOWN FOR VARIOUS MODELS IN 2020

Model	lockdown 1 MAPE [%]	no-lockdown MAPE [%]	Relative Difference
ELES	3.55	2.68	32.46%
XLAB base long	4.31	2.65	62.64%
XLAB specday long	3.53	2.71	30.26%
XLAB base	3.95	2.91	35.74%
XLAB mobility	3.90	2.95	32.20%
XLAB mobility specday	3.65	2.94	24.15%

TABLE VI
MAPE FOR DIFFERENT MODELS UNDER VARIOUS LOCKDOWN
CONDITIONS WITH RELATIVE DIFFERENCES

Model	Metric	lockdown 1 MAPE [%]	lockdown 2 MAPE [%]	no-lockdown MAPE [%]
ELES	mean	3.55 (-17.63%)	2.38 (+3.48%)	2.68 (+1.13%)
	std	3.07	2.12	2.44
	max	16.42	15.45	30.66
XLAB base long	mean	4.31 (0.00%)	2.30 (0.00%)	2.65 (0.00%)
	std	3.38	2.09	2.48
	max	18.55	13.40	29.55
XLAB specday long	mean	3.53 (-18.10%)	2.38 (+3.48%)	2.71 (+2.26%)
	std	2.67	1.93	2.54
	max	16.10	12.28	27.53
XLAB base	mean	3.95 (-8.35%)	2.39 (+3.91%)	2.91 (+9.81%)
	std	3.14	2.07	2.90
	max	18.85	12.28	30.53
XLAB mobility	mean	3.90 (-9.51%)	2.70 (+17.39%)	2.95 (+11.32%)
	std	3.15	2.14	2.59
	max	18.24	14.91	27.94
XLAB mobility specday	mean	3.65 (-15.31%)	2.25 (-2.17%)	2.94 (+10.94%)
	std	2.81	1.78	2.64
	max	13.97	12.23	28.74

Their model performed impressively well for 2020, as we demonstrated in Table IV. Comparison to the ELES model is not relevant here, as we are focusing on assessing the impact of adding features in a controlled environment. As mentioned the only difference between the XLAB models are features.

With that in mind, when further observing Table VI, a pattern is observed demonstrating that models utilising either mobility or special day performed better compared to those not using it.

Another thing to notice is that the best performance was actually during the second lockdown, which can be observed over all models. This could be attributed to various facts, where patterns could have stabilised by that point and become more predictable. Additionally, models were updated with new consumption patterns by then.

The last observation from Table VI is that the mobility and special day features did not significantly enhance performance on a typical non-COVID-19 day. To confirm this observation, we will perform an extensive study of a variety of features in the next Section III-D.

D. Effect of Additional Features to Model Performance

In this section, we utilize linear regression, XGBoost, and CatBoost models to assess feature importance across a variety

of input features. The results were evaluated and averaged for the years 2020 to 2023, and are presented in Table VII.

The first experiment, "None", does not include any additional features besides the signal itself. The next two experiments contain the signal and date-time features, along with off-time features such as weekends and holidays. Together, these features form a base, which is used in subsequent experiments to study the impact of individual features.

The first subgroup of experiments includes base features and weather features. As shown in Table VII, adding future weather improves performance for more complex models, where experiments utilising base and future weather yielded the best overall results for XGBoost and CatBoost.

The next group of experiments examines the addition of mobility data to the base set of features. For linear regression models, the mobility group outperformed the weather features group, whereas, for more complex models (XGBoost and CatBoost), the performance was worse compared to using weather features. Overall, more complex models performed better than the simple linear regression model. Here, we have to keep in mind, that errors are a lot higher for linear regression.

TABLE VII
EFFECT OF ADDING NEW FEATURES BETWEEN 2020 AND 2023.

Experiment Name	linreg MAPE[%]	xgb MAPE[%]	cbm MAPE[%]
none	6.20 (+0.00%)	5.46 (+0.00%)	5.13 (+0.00%)
datetime	6.66 (+7.31%)	3.71 (-32.07%)	3.37 (-34.28%)
base	5.67 (-8.60%)	3.60 (-34.01%)	3.39 (-33.85%)
base_weather	5.44 (-12.35%)	3.69 (-32.36%)	3.36 (-34.43%)
base_weatherfut	5.63 (-9.22%)	3.51 (-35.67%)	3.17 (-38.29%)
base_mob	5.42 (-12.63%)	3.67 (-32.87%)	3.36 (-34.58%)
base_specday	5.64 (-9.01%)	3.64 (-33.40%)	3.36 (-34.49%)
base_mobspecday	5.31 (-14.38%)	3.67 (-32.80%)	3.31 (-35.45%)
base_mobfuture	6.94 (+11.95%)	3.67 (-32.86%)	3.31 (-35.52%)
all_specday	5.40 (-13.00%)	3.64 (-33.37%)	3.23 (-37.01%)
all_mob	5.38 (-13.19%)	3.63 (-33.49%)	3.23 (-37.06%)
all_mobfuture	7.31 (+17.79%)	3.63 (-33.47%)	3.23 (-37.03%)
all	6.06 (-2.27%)	3.64 (-33.42%)	3.19 (-37.73%)

*base stands for date-time and holiday features used together

In the final set of experiments, combining all features, including special days, mobility, and future mobility, resulted in overall solid performance. The complex models, XGBoost and CatBoost, showed consistent improvement and leveraged the extensive feature set effectively. This indicates that the integration of a diverse range of features allows these advanced models to extract and utilize information more effectively, enhancing their prediction accuracy. Notably, the models in the "all" group demonstrate that these models performed quite well on average, achieving significant reductions in error rates compared to the baseline models. Even though the best results were achieved using future weather data, it makes sense to include data related to mobility, if available.

In the next Section, we focus on the effect of given features during the COVID-19 lockdown period in 2020.

E. Effect of Additional Features on Model Performance during COVID-19

To evaluate the impact of additional features during the COVID-19 pandemic, we must analyze their effects specifically during lockdown periods. Let’s first examine Table VIII presenting a performance of the models during normal or non-lockdown days. In comparison to the findings in Table VII (from the previous Section), we observe a consistent pattern: models incorporating mobility and special day features offer limited performance improvements.

TABLE VIII
EFFECT OF ADDING NEW FEATURES FOR NON-LOCKDOWN DAYS IN 2020 (MAPE)

Experiment Name	linreg MAPE[%]	xgb MAPE[%]	cbm MAPE[%]
none	6.83 (0.00)	5.99 (0.00)	5.55 (0.00)
datetime	9.81 (43.70%)	4.22 (-29.54%)	3.68 (-33.82%)
base	6.94 (1.61%)	3.99 (-33.31%)	3.78 (-31.82%)
base_weather	6.25 (-8.43%)	4.14 (-30.87%)	3.69 (-33.52%)
base_weatherfut	6.78 (-0.68%)	3.93 (-34.37%)	3.50 (-37.00%)
base_specday	6.95 (1.90%)	3.99 (-33.31%)	3.76 (-32.36%)
base_mob	6.83 (0.14%)	4.32 (-27.84%)	3.94 (-29.07%)
base_mobfuture	12.16 (78.10%)	4.33 (-27.68%)	3.91 (-29.57%)
base_mobspecday	6.85 (0.32%)	4.32 (-27.86%)	3.93 (-29.29%)
all_specday	6.12 (-10.35%)	4.12 (-31.23%)	3.69 (-33.18%)
all_mob	6.77 (-0.81%)	4.25 (-28.93%)	3.75 (-32.38%)
all_mobfuture	12.59 (84.47%)	4.29 (-28.32%)	3.79 (-31.75%)
all	6.95 (1.89)	4.25 (-29.13%)	3.78 (-31.99%)

**base stands for date-time and holiday features used together*

We must keep in mind that approximately **68.2%** of our data represents normal, non-lockdown days, with the remainder being lockdown days. This context is crucial for interpreting model performance. Similar to our observations with TFT, we face the challenge of limited data (only one year), which may affect prediction accuracy. Nonetheless, patterns forming across many models and feature combinations should still yield relevant results to be able to confirm or not confirm our hypothesis.

Table IX confirms a significant drop in model performance during the first lockdown compared to normal days, aligning with the results from Table V. Furthermore, linear regression was the least accurate of the models, making it almost unuseful in some cases. While we cannot expect lockdown prediction accuracy to match non-lockdown periods, our goal should be to minimize this performance reduction.

What we can notice, is that experiments including mobility data and special day data do contribute to improvements to better performance across all models and combinations. Several factors support this observation. The first one is that all best-performing experiments for every model include mobility or special days. For the second one, let us focus on the CatBoost model. Based on the results, it is the best-performing model.

Before we analyse the results more in-depth. It is worth clarifying that only experiments incorporating special days and mobility futures contain information on possible big changes in consumption patterns. While the ‘mobility future’ scenario demonstrates potential gains with perfect mobility forecasts,

TABLE IX
EFFECT OF ADDING NEW FEATURES FOR THE FIRST LOCKDOWN DAYS IN 2020 (MAPE)

Experiment Name	linreg MAPE[%]	xgb MAPE[%]	cbm MAPE[%]
none	8.33 (0.00)	8.19 (0.00)	7.57 (0.00)
datetime	8.68 (4.17%)	6.67 (-22.61%)	5.86 (-18.56%)
base	8.20 (-1.58%)	6.98 (-14.81%)	6.78 (-10.44%)
base_weather	7.43 (-10.81%)	7.25 (-11.58%)	6.99 (-7.64%)
base_weatherfut	8.41 (1.06%)	6.99 (-14.75%)	6.36 (-15.93%)
base_specday	6.98 (-16.21%)	6.98 (-14.82%)	6.17 (-18.50%)
base_mob	8.34 (0.19%)	7.03 (-14.15%)	6.40 (-15.48%)
base_mobspecday	6.92 (-16.90%)	7.03 (-14.15%)	5.64 (-25.53%)
base_mobfuture	12.23 (46.89%)	6.99 (-14.75%)	5.49 (-27.40%)
all_specday	6.14 (-26.25%)	7.30 (-10.97%)	5.78 (-23.66%)
all_mob	8.29 (-0.44%)	7.21 (-11.97%)	6.68 (-11.80%)
all_mobfuture	14.49 (73.93%)	7.18 (-12.38%)	6.65 (-12.12%)
all	13.47 (61.70%)	7.21 (-11.97%)	5.95 (-21.34%)

**base stands for date-time and holiday features used together*

this is not realistically achievable. For the CatBoost model best-performing experiment was the ‘mobility future’, since it is not realistically possible, it is presented in italic.

Overall, we can notice that CatBoost models containing either mobility, special day or both performed better compared to those not utilising these features. However, it’s important to stress that this advantage is not observed for non-lockdown days.

When analysing the effect of individual features, based on the results we could argue that special day feature has a bigger impact than mobility. Of course, in a world where we would be able to perfectly predict mobility, the best feature would be (future) mobility.

Across the tables VII, IX, and VIII, we observe a pattern: models containing all features often under-perform compared to the best-performing combinations. This suggests the potential impact of the ‘curse of dimensionality’. With too many features and limited data, the model may struggle to identify meaningful relationships.

More importantly, the pattern observed is similar to that of Section III-C, where we have demonstrated that models utilising either mobility or special day feature, on average performed much better during the first COVID-19 lockdown, compared to those not utilising. Similar conclusions can be made on effect during normal days, where their effect is present but does not have a big impact, in some cases even causing a curse of dimensionality.

F. Explainability and Feature Importance

To address the explainability part of the model, we provide average feature importance from the iterative model for the past three years as seen in Figure 2. The encoder features represent information from the past that is already known, while the decoder features represent information from the future that we are trying to predict.

The encoder features (in blue) in Figure 2 demonstrate the dominance of real load in shaping end results within the encoder. The second most important feature is mobility, which additionally confirms our observations. Date-time features and




Fig. 2. Encoder feature importance

modelled solar radiation follow in importance. Individual modelled weather features exhibit relatively minor impact; however, their aggregate influence remains significant due to their large number.

The decoder features (in orange) in the same Figure 2 reveal a similar trend to encoder features, particularly regarding the most influential factors. Notably, modelled weather features appear to exert a stronger impact within the decoder context.

While it is expected for target and date-time features to be most important, it was less expected to see a big impact of measured solar irradiation on predictions. This could be explained by an increasing amount of PV installed. In Slovenia, solar energy accounts for less than 10% of total energy (more precisely 7.58% as of 2023 [22]). This growth in solar power could be making accurate solar irradiation data even more important for accurate predictions.

IV. DISCUSSION

The obtained results show, that mobility data and special day features do improve the results during the first lockdown period. This was demonstrated for deep learning models utilising an iterative approach, as well as gradient boosting methods such as XGBoost and CatBoost. In the case of the iterative approach, the performance gains were almost 20 % for the model containing special days only and the model with no context regarding the COVID-19 status, supporting the findings in [3].

On the other hand, we have demonstrated the opposite for normal days, non-COVID-19 days, where the addition of mobility and special day features did not significantly improve the results. Even though no significant improvement was present, feature importance for multiple runs in 2020 suggests that mobility contains a lot of relevant information being ranked the second most important in the feature importance plot in Figure 2. The information it provides might be extracted elsewhere from weather and energy consumption patterns. The overall TFT model, when applied to load demand prediction, outperforms current state-of-the-art approaches. We have demonstrated that the iterative model is able to ingest new consumption patterns, thus improving its performance.

These results verify the hypothesis proposed in the context of COVID-19 pandemic data, demonstrating that model accuracy and performance can be improved during large-scale events. To address the more subjective research question of what energy critical infrastructure (CI) should consider in disruptive scenarios, recent experiences and best practices from the pandemic indicate that the primary focus should be on understanding the impact of demand changes on a transmission system (TS). It is essential for a CI operator, specifically a TSO, to always be prepared for any unforeseen scenarios, such as future pandemics, to maintain operational continuity (flawless, uninterrupted core business activities). At the same time, it is crucial to uphold activities related to grid resilience, employee health and safety, cybersecurity measures, and other vital operations.

The severe disruptions caused by the COVID-19 pandemic in the daily routines and behaviours of consumers led to a change in the shape of the daily consumption diagram. For instance, before COVID-19 the daily peak of demand was in the evening, however during COVID-19 the daily peak shifts to the morning. The previous change has a big influence on the adequacy assessment of a TS. Namely, the adequacy analysis or adequacy assessment of a TS is the most important analysis of a TS which answers the question: "Is there going to be enough energy in the system in each situation including failures of generators". Based on this analysis, the levels of necessary spinning and non-spinning reserves are determined which are crucial for the secure operation of a TS. An incorrect adequacy analysis leads to an increase in the percentage of loss of load probability (LOLP), which can further lead to significant economic damage. The economic damage when the energy is not served (ENS) is usually estimated with the Value

of Lost Load (VoLL).

VoLL represents the economic value associated with not being able to supply electricity to consumers during periods of high demand or supply shortages. It reflects the cost to consumers of being without electricity for a certain period and includes factors such as lost production, inconvenience, and potential damage to equipment or goods. VoLL is usually expressed in currency per MWh and can vary depending on factors such as the type of consumer, the time of day, and the duration of the outage. In the European Union, the specific values for VoLL can vary between countries and regions due to differences in electricity market structures, consumer preferences, and economic conditions. Typically, VoLL values in the EU range from around €1,500 to €23,000 per MWh [23], but they can be higher in some cases, particularly for critical services or industries where the cost of downtime is very high. For comparison, wholesale market prices are usually at the level between 50-100 €/MWh. Other key aspects to be considered are discussed in [24].

The conventional approach takes typical daily load patterns when conducting an adequacy analysis. The results of this research hint at the issues that energy CIs must consider when preparing resilience strategies for disruptive scenarios based on recent experience and good practices in the latest pandemic. They show that the typical daily load patterns can be interrupted for a significant period in the case of large-scale events such as COVID-19. Therefore, to decrease LOLP the CI (i.e. TSOs) should consider the use of untypically daily load patterns while conducting an adequacy analysis. These load patterns can be generated by the model described in this research.

Regarding further analysis and discussions for future disruptive scenarios, TSOs could also consider using industrial and residential power consumption predictions as encoder and decoder features (for total consumption forecast), taking into account that these two consumption categories could perform significantly differently in specific situations. Currently, most of the TSOs do not have separate metering information for industrial and residential power consumption in real-time (usually they get data for residential consumption at the end of the month), as well as their forecasts, although these values could be estimated well.

Also, the number of prosumers or active customers (solar installations "behind the meter") is increasing rapidly, and taking into account that analysis shows solar irradiation as a very important feature, future models could also consider the installed power capacity of active customers/prosumers (at least on yearly level) as a feature (because it is changing/growing over the time, and solar generation has the impact on the measured net consumption).

V. CONCLUSIONS

The goal was to develop a resilient model, that enables us to make better predictions during times of large-scale events that have a significant effect on energy demand prediction models. In this work we show how we can use the mobility and

lockdown flag to improve model accuracy and performance during impact large-scale events, taking as a basis the data collected during the COVID-19 pandemic.

We achieve these improvements with additional features, particularly by considering special day features. The biggest impact of these features can be observed in the first lockdown period, whereas for normal days, improvements were harder to notice. Overall it makes sense to utilise both mobility and special day functions, if available. They offer insights in case of large-scale events, even if they are not common. When taking into account the cost of individual features, mobility data may turn out to be relatively expensive, here special day features are much more cost-efficient.

Moreover, the periodic monthly update of the model shows great benefit for the predictions computed and the resilience of the models, as we have seen that the effect of COVID-19 was hardly noticeable in the second lockdown.

It was particularly clear, in the cases of EKC and ELES, that the improvements in prediction accuracy have significant real-world benefits. During the COVID-19 pandemic, shifts in daily consumption patterns, such as peak demand moving from evening to morning, impacted the adequacy assessment of transmission systems (TS). Accurate forecasts help ensure there is enough energy to meet demand, reducing the LOLP and preventing economic damage from unserved energy (ENS). Better predictions lead to optimal levels of spinning and non-spinning reserves, enhancing grid stability. Economically, accurate predictions minimize wasted energy, saving costs associated with VoLL.

The investigation of larger, more complex versions of the TFT is crucial as the industry shifts from 1-hour to 15-minute resolution quadrupling the input parameters. We have demonstrated that using a 168h input window improves its performance, but again increases the number of input features by up to seven-fold. Both changes result in a much bigger and more complex model, highlighting a key area for future research on performance impacts.

Research is increasingly focusing on foundational time-series models based on transformers. Authors of research work in [25], [26], [27] and [28] focus on zero-shot forecasting of univariate time-series. Inspired by breakthroughs in natural language processing with models like Gemini, GPT and Claude, time-series prediction, similarly aims to predict the most probable next value based on prior input. Referenced models often outperform TFT in certain applications, showing promising results. These methods do not directly apply to our work, as we are using multivariate time-series, whereas the papers are focusing on univariate time-series. A potential implementation approach for multivariate foundational models could be composed out of multiple foundational models, each fine-tuned for a specific task, then fusing the outputs to create a similar architecture to the TFT. At the time of writing, no studies have been published on this specific approach, but further research in this direction is anticipated. Overall, these approaches will play a significant role in efficiently managing the operations of CIs.

ACKNOWLEDGMENT

This research has received funding from the European Union's Horizon Europe research and innovation programme under the grant agreement No. 101073821 (SUNRISE) and 101070052 (TANGO).

REFERENCES

- [1] M. Kwiatkowska and X. Zhang, "When to trust ai: Advances and challenges for certification of neural networks," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35. IEEE, 2023. doi: 10.15439/2023F2324 pp. 25—37.
- [2] S. D. Li, C. and T. Reindl, "Real-time scheduling of time-shiftable loads in smart grid with dynamic pricing and photovoltaic power generation," in: *2015 IEEE Innovative Smart Grid Technologies - Asia (ISGT ASIA)*, pp. 1–6, 2015.
- [3] Y. W. Chen, Y. and B. Zhang, "Using mobility for electrical load forecasting during the covid-19 pandemic," 2020.
- [4] B. Lim, S. Arnk, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021.
- [5] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," *CoRR*, vol. abs/1603.02754, 2016. [Online]. Available: <http://arxiv.org/abs/1603.02754>
- [6] A. V. Dorogush, V. Ershov, and A. Gulin, "Catboost: gradient boosting with categorical features support," 2018.
- [7] J. Jenko, "Development and analysis of new activation based load profiles," Master's thesis, University of Ljubljana, Faculty of Electrical Engineering, 2023.
- [8] J. Jakob and J. Pita Costa, "Using temporal fusion transformer predictions to maximise use of renewable energy sources," in: *Proceedings of the International Workshop on Artificial Intelligence and Machine Learning for Energy Transformation (AIE)*, IEEE, 2024.
- [9] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," 2020.
- [10] Optuna. (2024) Optuna documentation. [Online]. Available: <https://optuna.readthedocs.io/en/stable/>
- [11] Pytorch. (2024) Quantileloss documentation. [Online]. Available: https://pytorch-forecasting.readthedocs.io/en/stable/api/pytorch_forecasting.metrics.quantile.QuantileLoss.html
- [12] Meteostat. (2024) Weather and climate database. [Online]. Available: <https://meteostat.net/en/>
- [13] Open-meteo. (2024) Free weather api. [Online]. Available: <https://open-meteo.com>
- [14] Openweathermap. (2024) Open weather history forecast bulk. [Online]. Available: <https://openweathermap.org/api/history-forecast-bulk>
- [15] ENTSO-E. (2024) Entso-e transparency platform. [Online]. Available: <https://transparency.entsoe.eu/>
- [16] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Ro, M. Wiebe, P. Peterson, P. Gerard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with numpy," *Nature*, vol. 585, no. 7825, p. 357 to 362, 2020. doi: 10.1038/s41586-020-2649-2. [Online]. Available: <https://doi.org/10.1038/s41586-020-2649-2>
- [17] W. McKinney *et al.*, "Data structures for statistical computing in python," in *Proceedings of the 9th Python in Science Conference*, vol. 445. Austin, TX, 2010, p. 51 to 56.
- [18] J. D. Hunter, "Matplotlib: A 2d graphics environment," *Computing in Science and Engineering*, vol. 9, no. 3, p. 90 to 95, 2007. doi: 10.1109/MCSE.2007.55
- [19] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, p. 261 to 272, 2020. doi: 10.1038/s41592-019-0686-2
- [20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An open source machine learning framework," *arXiv preprint arXiv:1802.02611*, 2018.
- [21] PyTorch Forecasting Developers, "Pytorch forecasting documentation," <https://pytorch-forecasting.readthedocs.io/en/stable/getting-started.html>, 2024.
- [22] P. Jowett. (2024) Slovenia's new solar additions hit 400 mw in 2023. [Online]. Available: <https://www.pv-magazine.com/2024/02/19/slovenias-2023-solar-additions-hit-400-mw/>
- [23] Swinand, G. Peter, Natraj, and Ashwini, "The value of lost load (voll) in european electricity markets: Uses, methodologies, future directions," in *2019 16th International Conference on the European Energy Market (EEM)*, 2019. doi: 10.1109/EEM.2019.8916400 pp. 1–6.
- [24] D. R. M. L. Quang Hieu Vu, Ling Cen, "Key factors to consider when predicting the costs of forwarding contracts," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, vol. 30. IEEE, 2022. doi: 10.15439/2022F293 pp. 447–450.
- [25] K. Rasul, A. Ashok, A. R. Williams, H. Ghonia, R. Bhagwatkar, A. Khorasani, M. J. D. Bayazi, G. Adamopoulos, R. Riachi, N. Hassen, M. Biloš, S. Garg, A. Schneider, N. Chapados, A. Drouin, V. Zantedeschi, Y. Nevmyvaka, and I. Rish, "Lag-llama: Towards foundation models for probabilistic time series forecasting," 2024.
- [26] A. F. Ansari, L. Stella, C. Turkmen, X. Zhang, P. Mercado, H. Shen, O. Shchur, S. S. Rangapuram, S. P. Arango, S. Kapoor, J. Zschiegner, D. C. Maddix, H. Wang, M. W. Mahoney, K. Torkkola, A. G. Wilson, M. Bohlke-Schneider, and Y. Wang, "Chronos: Learning the language of time series," 2024.
- [27] A. Das, W. Kong, R. Sen, and Y. Zhou, "A decoder-only foundation model for time-series forecasting," 2024.
- [28] A. Garza, C. Challu, and M. Mergenthaler-Canseco, "Timegpt-1," 2024.

HPC operation with time-dependent cluster-wide power capping

Alexander Kammeyer*[†] [0000–0002–7858–0354], Florian Burger* [0000–0003–4745–5515],
Daniel Lübbert* [0000–0003–3852–5665] and Katinka Wolter[†] [0000–0002–8630–0869]

*Physikalisch-Technische Bundesanstalt, Abbestraße 2-12, 10587 Berlin, Germany

Email: {alexander.kammeyer, florian.burger, daniel.luebbert}@ptb.de

[†]Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany

Email: {a.kammeyer, katinka.wolter}@fu-berlin.de

Abstract—HPC systems have increased in size and power consumption. This has led to a shift from a pure performance centric standpoint to power and energy aware scheduling and management considerations for HPC. This trend was further accelerated by rising energy prices and the energy crisis that began in 2022.

Digital Twins have become valuable tools that enable energy and power aware scheduling of HPC clusters. This paper uses an existing Digital Twin and extends it with a node energy model that allows the prediction of the cluster power consumption. The Digital Twin is then used to simulate system-wide power capping for different energy shortages functions of varying degree. Different policies are proposed and tested towards their effectiveness in improving the job wait times and overall throughput under limiting conditions.

Based on a real world HPC cluster, these policies are implemented. Depending on the pattern of the energy limitation and workload, improvements of up to 40 percent are possible compared to scheduling without policies for these conditions.

I. INTRODUCTION

HIGH-Performance Computing (HPC) is used, when a single computer is either too slow or too small for a single problem. Multiple computers, so called compute nodes or just nodes, are connected to solve the problem cooperatively. HPC systems have grown in size and capability over the years [1]. Simultaneously, their energy consumption did also grow, with the current top systems using exceeding 20 MW. Energy prices have increased as well, especially since the Russian war against Ukraine and the energy crisis in Europe in 2022 and 2023. The Physikalisch-Technische Bundesanstalt (PTB) operates a HPC cluster with approximately 30 kW installed power for research purposes.

During the 2022 energy crisis [2], Germany implemented "Ordinances on energy saving" [3] that contained measures to reduce energy consumption together with an appeal to public institutions, companies and private households to reduce their overall energy usage. In the case of an immediate energy shortage, power capping and load shedding measures have been discussed. These measures would require immediate power reduction on short notice. As part of PTB's strategy [4] towards energy-efficient HPC, continuous cluster operation under reduced energy availability is one of the goals. This paper presents policies and scheduling strategies for an HPC cluster to continue operation under these conditions.

Digital Twins are virtual representations of real-world objects, such as HPC systems. They collect data about the object and contain models that allow to simulate the behaviour and states of the real object. Over the past decade, they have grown in popularity in many industry applications and begin to see adoption in the HPC domain as well. Scheduling simulations are common in the area of HPC system research and an ideal basis for an HPC Digital Twin as they allow to model the system. In this paper they are used to test and verify the policies.

Figure 1 shows the mean weekly power consumption of the PTB campus where the HPC cluster is located. It shows a base power consumption of the campus at around 270 kW with 5 peaks for the traditional five workdays with very minor peaks on the weekend. With a possible power limitation, there would still be a similar pattern with a day-night-cycle for the workdays. This would allow a cluster operation at night with possible fewer restrictions. While the cluster allows very long jobs, a shift to a day and night cycle requires restrictions on the job length, so that jobs can be completed at night.

The trivial solution to power limitations is to turn off the HPC cluster entirely. The campus energy management system can also forcefully disconnect the HPC cluster or other large consumers from the power supply, if required. This does not guarantee a graceful shutdown and might lead to data loss. However, depending on the limitations of the energy usage, the cluster might remain operational under reduced load or with some of the nodes turned off. The system administrator could turn off nodes manually. This is a very coarse-grained approach. With a more fine grained approach towards power capping of the system and knowledge of the node power consumption, more nodes can remain online. Another common technique is Dynamic Voltage and Frequency Scaling (DVFS). With DVFS, the speed of the processors and thus the energy consumption of the compute nodes can be set dynamically. The Digital Twin of the HPC cluster can be used to estimate the energy consumption of individual jobs and plan accordingly while monitoring the overall system power usage to guarantee the operation within the defined limits. This entire process is also automatic, requiring no manual intervention from the operators.

This paper makes the following contributions:




Figure 1: Mean power consumption on the west-side of PTB's Berlin site. The HPC cluster is located in this part of the campus. The graph shows the mean weekly power consumption of 2023 in 15 minute intervals.

- Support for heterogeneous HPC systems with different nodes is implemented for the scheduling simulation in the Digital Twin. Previously, the simulation could only handle one node type. The simulation can also simulate DVFS of the individual nodes with regard to runtime and power consumption. For this purpose, a node energy model is created.
- Different policies, including scheduling fewer jobs to decrease overall system load, shutting down unused compute nodes and DVFS on a per-job level with the node energy model, are implemented in the scheduling simulation.
- The Digital Twin is extended with these policies and through scheduling simulations with different job traces and power limits it is tested that the Digital Twin ensures the system-wide power cap of the cluster. This allows continuous cluster operation under reduced energy availability without manual intervention.

II. RELATED WORK

With increasing energy consumption of HPC systems, research focuses more and more on energy and power related questions. The survey by [5] gives an overview of tools for energy and power management in contemporary HPC systems from a single node all the way to grid systems. It also shows research towards power prediction for different parts of HPC systems. A similar survey by [6] is done about power-aware scheduling.

The energy management framework for supercomputers (EAR) [7] is an accounting, control and optimisation framework for HPC systems. This system requires MPI profiling to create job profiles and uses an algorithm called DynAIS to detect different hotspots within an application. An Energy-

aware job scheduling strategy on top of EAR [8] tries to place jobs in heterogeneous clusters based on the job profile.

Similar to the node energy model in chapter III-D, a power profile for different applications or types of applications have been done for WZ factorisation [9] and matrix factorisation [10]. They also run a benchmark on different frequencies but did not determine the pareto-optimal frequencies. The results cannot be used directly here because of different hardware architectures and a different selection of applications. While the node energy model and the papers look at the power cap from the frequency side, the power cap can directly implemented via a true power limit through a driver yielding similar results [11].

A Digital Twin, as defined by [12], is a virtual representation of a real-world object, in the scope of this paper an HPC cluster. Several data inputs are integrated into the Digital Twin, which handles the incoming data, processes it and stores it. With a bi-directional link between the real-world object and the Digital Twin, they can influence each other as changes in the real world are reflected in the virtual world and vice versa. This requires regular synchronisation to ensure consistency, however, by definition, a permanent, immediate synchronisation is not necessary. The Digital Twin must enable interoperability with other systems. The concept of the Digital Twin has gained such importance and has found wide adoption, that it has been standardised by ISO [13].

Digital Twins of HPC systems aid the system operators by allowing to test configuration changes, policies and different scheduling algorithms with altering the actual cluster. Possible negative effects are avoided this way. Simulating the scheduler is common practice in HPC research. Different simulations have been created, e.g. based on the Slurm scheduler [14], [15], [16] or based on Digital Twins [17]. Scheduling in regard to power consumption and pricing has been demonstrated by [18] for different billing strategies by delaying jobs for a static price model. A day and night price model with a 0-1-knapsack strategy is presented by [19]. With the transformation to renewable energies in the energy production, a scheduler can also take CO_2 emissions into account as shown by [20].

III. PRELIMINARY WORK

This section presents the Digital Twin for the HPC cluster of PTB and how it has been extended for this paper with support for heterogeneous nodes and a power model. The power model is later used for the power prediction for the power capping.

A. PTB's HPC cluster

The PTB operates a relatively small cluster at around 30 kW installed power. This amounts to approximately 10 percent of the overall power consumption on campus (Figure 1). The cluster is equipped with two different CPU nodes. Nodes of the first node type are equipped with two Intel Xeon E5-2690 v4 [21] each. The second node type is also a dual-socket system with two Intel Xeon Gold 6132 CPUs [22] per node. Both processor types are operating at a base frequency




Figure 2: Network overview of all Digital Twin data sources, agents and network boundaries. [20]

of 2.6 GHz. Each node of the first type has 256 GB of DDR4 RAM while each node of the second type has 192 GB. The CPU can operate at frequencies from 1.2 GHz to 3.5 GHz (1.0 GHz to 2.6 GHz), for the first (second) node type, respectively. The manufacturer specified turbo frequencies of up to 3.7 GHz are disabled on the second node type. The cluster provides a total of 60 nodes for applications with 24 nodes of the first type and 36 nodes of the second type.

Both CPU designs are dated. The first was released in 2016 while the second was launched one and a half years later in 2017. Newer CPU generations can be expected to be more energy efficient. In fact, Moore’s Law [23] states that the number of transistors doubles every one and a half years. Closely linked to the number of transistors is the performance and often also energy consumption. For the purpose of this study, the performance of the two designs are used as examples. The mechanisms presented in later chapters work independent of specific CPU designs and the underlying hardware is configured through the node energy model, presented in section III-D, in the Digital Twin.

B. Digital Twin Components and Layout

PTB is actively developing a Digital Twin for their HPC cluster [20]. The Digital Twin integrates all sensor data of the data centre as well as the cluster itself. Most of the data is time-series measurement data, thus the central database of the Digital Twin is an InfluxDB. The data centre is equipped with a multitude of sensors, such as electric energy meters, heat flow meters and temperature sensors. Additionally, external data sources, such as energy generation, weather forecast and energy CO₂ intensity, are also integrated into the Digital Twin. For security reasons, the sensors are separated into a so-called building infrastructure network. The HPC cluster also has a separate network. The Digital Twin exchanges data via well-defined interfaces between these networks. The entire network layout with all sensors, meters and data sources is shown in Figure 2. The Digital Twin also contains a scheduling simulation as the digital representation of the cluster behaviour. This simulation uses the data from the InfluxDB to get the system state and can use job traces to test system configurations. This allows the system administrator to test configurations with altering the production system. Simulations are also cheaper,




Figure 3: Collector agent schematic with 3-stages for data collection, transformation and output to the database. [20]

since they do not require actual jobs to run on the system and thus use far less energy. The simulation is event-based and use job traces from the cluster or use the parallel workload format [24].

C. Data Collector Agent

The collector agents use a three-stage pipeline to collect, transform and send data to the central database. The sensors and data sources each use different protocols such as REST/JSON, XML and M-BUS. The database uses a custom format that can be accessed through a client library. The agent first collects the data through the appropriate protocol, then transforms the data into a format supported by the database and finally sends it to the database. This allows to add new sensors and protocols easily, as only the corresponding stages need to be adapted. The architecture of the agent is shown in Figure 3.

D. Node Energy Model

The cluster has different nodes with different hardware and thus different energy consumption. A node energy model has been created that allows the simulation to trace the energy consumption of different job types running at different frequencies. Other frameworks, such as EAR [7], rely on injecting code into the application to trace job behaviour. For this study, a less intrusive approach has been chosen. Two benchmarks and two applications have been selected and run at different frequencies to create the node energy model. The selected applications represent common applications on the HPC cluster.

The High-Performance Linpack (HPL) [25] is a common HPC benchmark and is also used to create the TOP500 ranking. It is a highly optimised linear equation solver. The benchmark uses vector instructions like AVX. These instructions require large amounts of energy and thus the benchmark is suitable as an upper bound for the node energy model. It is unlikely that a real application uses more energy than the HPL benchmark.

Another common HPC benchmark is the High Performance Conjugate Gradients (HPCG) [26]. The HPL benchmark is not representative of all HPC applications and the HPCG tries to complement this with a broader set of operations with different data access patterns that are harder to optimise than the pattern used by the HPL benchmark.

The first of the real applications is Open Field Operation And Manipulation (OpenFOAM) [27]. It is a computational

Table I: The Node Energy Model

Node Type	Job type	Frequency	Power	Scaling factor
1	HPL	3.2 GHz	415 W	0.966
		2.6 GHz	388 W	1.000
		2.2 GHz	322 W	1.163
	HPCG	3.2 GHz	147 W	0.884
		2.6 GHz	137 W	1.000
		1.5 GHz	108 W	0.932
	OpenFOAM	3.5 GHz	266 W	0.922
		2.6 GHz	240 W	1.000
		2.3 GHz	185 W	1.205
	Geant4	2.8 GHz	191 W	0.894
		2.6 GHz	183 W	1.000
		2.3 GHz	156 W	1.177
Idle	-	57 W	-	
Offline	-	5 W	-	
2	HPL	2.6 GHz	349 W	1.000
		2.3 GHz	349 W	0.992
		1.8 GHz	298 W	1.133
	HPCG	2.6 GHz	170 W	1.000
		2.3 GHz	161 W	0.926
		1.8 GHz	146 W	1.067
	OpenFOAM	2.6 GHz	273 W	1.000
		2.5 GHz	228 W	1.107
		2.3 GHz	212 W	1.188
	Geant4	2.6 GHz	219 W	1.000
		2.2 GHz	181 W	1.159
		2.0 GHz	167 W	1.311
Idle	-	51 W	-	
Offline	-	5 W	-	

fluid dynamics package that is commonly used by PTB researchers for investigating flows through pipes and other geometries.

The second selected application, Geometry and Tracking (Geant4) [28], is a Monte Carlo simulation toolkit for studying particles passing through matter. It is widely used in various fields such as high energy physics, medical physics and others.

Each of the four jobs was run on the two node types on all supported frequencies. The energy consumption of each job was monitored through IPMI. This allows to calculate two metrics for each frequency: Time-to-Solution (TtS) and Energy-to-Solution (EtS). Figures 4a and 4b show two such results for OpenFOAM and the HPL benchmark respectively on the newer node type 2. With regard to the two metrics TtS and EtS, the pareto front [29], [30] for each application and node type is computed, shown as the orange dots. These points are not dominated by any other point. For the purpose of this paper, from each optimal point set, two points have been chosen together with the processor base frequency of 2.6 GHz. This allows the simulation to choose from three points for each job type. Applying DVFS also changes the job execution time. The simulation adjusts for that by multiplying the job length with a factor based on the runtime of the job compared to the base frequency of 2.6 GHz. The node model also contains values for idle and offline power consumption of the nodes. A node consumes energy when offline, because the Wake-on-LAN functionality needs to listen for the magic packet. Table I summarises the frequencies, power and scaling factor.

IV. PROBLEM DEFINITION

With the begin of the energy crisis in Europe in 2022, rising energy prices and energy scarcity became a concern for HPC operators. The way a cluster uses energy can be controlled through the scheduler and resource manager. Possible scheduling policies in terms of energy cost, e.g. through delaying jobs, have been discussed in Chapter II and have been shown with a Digital Twin [4]. Regarding energy scarcity, a HPC cluster has a variable energy consumption mostly defined by the compute nodes and the jobs running on them. This paper focuses on policies that can be implemented to continue stable operation under reduced energy availability conditions when the limitations are known in advance.

Energy scarcity can arise through different factors: an insufficient power supply to the entire campus from the energy supplier or an insufficient distribution inside the campus resulting in a scarcity for the cluster. The power capping functionality can also be used to reduce the thermal output of the cluster in case of problems or limitations of the cooling infrastructure.


A data centre can be viewed as four pillars [31]: building infrastructure, system hardware, operating software and user applications. Changes to the first and second pillar are possible but often involve installing new hardware or larger construction work. Both is time consuming and not suitable for short term measures against shortages. This paper focuses on changes in the system software, the third pillar, that can be applied by the HPC system administrator immediately.

A. Proposed Policies

The main task of a batch scheduler is to allocate resources to jobs. The current cluster configuration requires users to select one of the two node types. Therefore, the scheduler must put jobs entirely on one of the two node types. The scheduler cannot select a different node type or make mixed allocations. If this criterion was to be relaxed, the scheduler could select nodes first, that have a lower EtS for the given job type.

The nodes contribute the most to the power consumption of the cluster. The energy model in Table I shows, that both types have an idle consumption of 57 W and 51 W respectively. This results in an idle consumption of 3204 W for the entire cluster. As an energy saving measure, nodes that are currently not used, can or must be turned off to stay below the power limit. An interval of 5 minutes has been selected as policy for this case. If a node is idle longer, the Digital Twin will shut it down and restart it, if needed. The boot time of a node is about 2 minutes, after which the node can receive new jobs. If the power limits allows it, the Digital Twin will keep a certain number of nodes online so that new jobs have a chance to start immediately.

DVFS allows to scale the energy consumption of the nodes according to the current limitations. Figure 4a shows all frequencies of a node type. While the node uses less power on these lower frequencies, the EtS and TtS is much higher compared to frequencies from the middle of the frequency range, with the two lowest frequencies, 1.1 GHz and 1.0 GHz



(a) Pareto front for OpenFOAM on node type 2 with all supported frequencies. (b) Pareto front for HPL on node type 2. Frequencies below 1.5 GHz are not shown.

Figure 4: Two exemplary Pareto-fronts. Pareto-optimal frequencies are coloured orange. The remaining frequencies are coloured blue.

exceeding the highest frequency 2.6 GHz in terms of EtS. The decisions has been made, not to use these frequencies as they are too inefficient. For each job type and each node type, a set of frequencies has been select as describes in Chapter III-D.

Changing the frequency, and thus speed of the processors, implies a change in job runtime. If the scheduler forces a certain frequency, it adjusts the job runtime with the scaling factor from Table I. As described above, jobs run exclusively on one of the node types. If they ran on different types simultaneously, the simulation would have to adjust the overall runtime accordingly and respect the different speeds of the nodes. Since this is currently not enabled on the cluster, it has not been integrated in the simulation.

For the purpose of this simulation, all power limitations are known in advance. Currently the cluster allows a very high job length of up to four weeks. This makes it impossible to react to any new limitations on short notice and jobs need to be cancelled. Given the cyclic nature of the power consumption on campus (Figure 1), a shorter maximum job runtime is required. The issue with crisis induced power limits is that they may also be only known a few days or even just hours in advance. Although no concrete time frame has been set, the notice will most likely not come four weeks beforehand.

The simulation currently does not support the cancellation of jobs for power capping. The Digital Twin tries to start a job only if the energy limit allows the jobs and nodes to run. In a real-world scenario, jobs might be cancelled, e.g. if their power profile drastically exceeds previously observed profiles by the Digital Twin. Simply re-scheduling the job might not be feasible because a job is not guaranteed to be free from side-effects. In case of job cancellation, manual intervention by the user is necessary.

B. Algorithm

Algorithm 1 Pseudo-code of the core scheduling routine

```

1  backfill_power(eligible_jobs) {
2    for (Job j : eligible_jobs) {
3
4      // test if nodes need to be booted
5      // check if enough nodes are offline
6      if (j.nodes > online_nodes(j) &&
7          j.nodes <= online_and_offline_nodes(j)) {
8        // how many nodes need to be booted?
9        toboot = j.nodes - online_nodes(j);
10
11       // check if the nodes would
12       // exceed the power limit
13       if (!check_power(j, get_offline(j, toboot))) {
14         continue;
15       }
16       boot_nodes(j, toboot);
17     }
18
19     // test of enough nodes are online
20     // and available
21     if (j.nodes <= online_nodes(j)) {
22
23       // check if the job and nodes would
24       // exceed the power limit
25       if (!check_power_dvfs(j)) {
26         continue;
27       }
28
29       assign_nodes(j);
30       j.wait_time = tick - j.submit_time;
31       running.add(j);
32       eligible_jobs.remove(j);
33
34       // trigger an event in the simulation
35       // on job completion
36       e = new JobEvent(tick + j.run_time_scaled,
37                       j, JobState.COMPLETED);
38       eventQueue.add(e);
39     }
40   }
41 }

```

With the considerations from the previous section, a policy has been developed, that handles the cluster operation under reduced energy availability conditions. The scheduling simulation is event driven. Three different events are defined: a job event is triggered when a job is submitted or finished, a node event is triggered when a compute node gets assigned a job, finishes a job, starts or shuts down and a power event is triggered when the current power limit for the cluster changes. The simulation then handles the event and calls the scheduler. The core scheduling logic is shown as pseudo-code in Algorithm 1.

For each queued job, the algorithm first checks, if nodes need to be turned on (lines 6 and 7). *online_nodes()* returns the idle nodes the job can use and *online_and_offline_nodes()* checks whether enough nodes are idle or online. It does not make sense to boot nodes when the job cannot start. If that is the case, it is checked in line 13 with *check_power* whether the additional nodes, together with the projected energy consumption of the job, would exceed the power limit. The function *get_offline()* returns a list of offline nodes the scheduler would allocate to the job. If the limit is exceeded, the next job is checked. If not, the required nodes are booted. Each nodes takes two minutes to boot.

If a job has enough nodes available (line 21), the algorithm checks if the job fits within the limit with the *check_power_dvfs()* function in line 25. The function tests all DVFS settings from Table I and adjusts the job length if necessary. If the job power requirement exceeds the limit, the algorithm moves to next eligible job. Otherwise, the job gets assigned compute nodes, is started and removed from the eligible job list (lines 29-31). An event is created when the job finishes so that the scheduling simulation can mark the nodes available again. The event queue is a priority queue that sorts the events by their tick.

This algorithm is First Come, First Serve (FCFS) with backfilling. In this implementation, jobs are allowed to push back larger jobs. This design decision has been made because energy availability during a crisis situation is unclear. The cluster could stop operation entirely. Therefore, completing jobs gets precedence over the fairness criterion.

V. EVALUATION

The previous chapter introduced the proposed policies. This chapter focuses on their implementation in the Digital Twin. The current cluster allows long job run times. The average power usage on campus peaks at around 400 kW (Figure 1). For this experiment, a limit of 300 kW is assumed. The power limit for the cluster is the value between the average consumption and 300 kW. The limit is capped at 5 kW power usage for the cluster under the assumption that some part of the power budget gets allocated to the cluster. As shown in Figure 5a, a job trace from the real cluster together with this power limit only allows a few jobs to start in between the limits and some more during the weekend. The rest is kept in queue until the limit ends due to their size.

Table II: Results of the scheduling experiment with the two power limit patterns and averages for 5 different job traces each.


pattern	policy	wait time	sim run time
synthetic	limit	3.766 h	14.082 d
	shutdown	2.472 h	14.059 d
	dvfs	2.575 h	13.417 d
	dvfs + shutdown	2.289 h	13.414 d
real	limit	5.423 h	15.110 d
	shutdown	3.649 h	14.757 d
	dvfs	5.124 h	14.948 d
	dvfs + shutdown	3.415 h	14.663 d

As discussed in Chapter IV-A, shorter jobs can circumvent this issue. To generate such a job trace, the Feitelson job model [32], [33] is used. This model allows to configure the maximum job length together with some other parameters such as arrival rate of the jobs. It does not support different job types. They are generated at random and with equal probability using a discrete uniform distribution. For this experiment, two different limits are used for testing. The first one is the same as in Figure 5a, based on the campus power consumption (*real*). The second is a simpler, pyramid-shaped limit that uses the same lower bound of 5 kW (*synthetic*). Five different job traces with different seeds were generated for the following experiments.

The first step is to validate that the scheduler does not exceed the power cap. All compute nodes remained online and all jobs ran at the highest frequency. The scheduler can, if required, turn off individual nodes if it would otherwise exceed the power limit. They are re-booted when needed by a job. In a second step, the scheduler will turn off nodes proactively in order to save energy and thus allow more jobs to run. An idle node is turned off after 5 minutes. Finally, DVFS was added to the simulation. This allows the scheduler to start jobs at a lower frequency but also with less power. The results can be found in Table II.

The table contains two metrics: the mean job wait time and the overall run time of the simulation (Figures 5c, 6a and 6c). For the synthetic limit, the pro-active shutdown of compute nodes brings down the mean job wait time by 34 percent from 3.766 h to 2.472 h. Enabling DVFS gives an improvement of 32 percent. Combining node shutdown and DVFS further improves the wait time down to 2.289 h or 40 percent compared to the run without. The overall simulation run time on the other hand benefits more from DVFS than node shutdown. Here, the scheduler can start jobs at lower frequency and thus fit more jobs below the limit. This policy prefers smaller jobs which can in turn lead to an increase of the wait time for larger jobs.

The second experiment with the power limit based on the campus energy usage showed an improvement of 37 percent. The effect of DVFS on the wait time is smaller because of the longer periods of low power limits in this experiment. Combining both policies also yields a 5 percent improvement as in the first test case. For this simulation, the simulation length benefited more from the node shutdown than DVFS.



(a) Power trace of the PTB job trace showing four weeks of power limitations based on the campus energy usage. (b) Power trace of a Feitelson job trace showing two weeks of power limitations based on the campus energy usage. (c) Power trace of a Feitelson job trace showing three weeks of power limitations based on the synthetic power limit function.

Figure 5: Results of the experiments

For all experiments should be noted, that the results are dependent on factors such as job length, the pattern of the limitation and overall utilisation of the cluster. In this simulation, the lowest limit was at 5 kW while the idle consumption of the cluster is 3.2 kW. This leaves only 1.8 kW for compute jobs.

VI. CONCLUSION AND FUTURE WORK

This paper presented an approach to handling power limitations in the energy supply for an HPC cluster with the aid of a Digital Twin. The Digital Twin is able to trace and predict the power consumption of the HPC cluster. This allows the Digital Twin to control the HPC cluster and keep the overall power consumption below a defined threshold. This power capping capability is required, if the energy consumption of the HPC system needs to be limited.

Support for heterogeneous HPC clusters was added to the Digital Twin. Previously, only homogeneous HPC clusters with a single node type were supported. The Digital Twin can handle clusters with multiple types of nodes, each with an individual power consumption. Based on an example HPC cluster at PTB, a node power model was created (Table I). This model allows the Digital Twin to use DVFS to run jobs at different clock speeds and influence their energy consumption.

This paper proposed three strategies to continue operation under reduced load: trace power consumption and only start jobs that fit in the power budget, shutdown idle nodes when they are idle for a certain amount of time and DVFS to allow jobs to start with a reduced frequency set by the scheduler and also reduced power consumption (Algorithm 1).

These strategies were then compared with two different energy limits: one synthetic energy limit and one energy limit based on the campus energy consumption. The simulation showed, that the cluster can stay below the power cap and the target metrics were improved with the proposed policies, in some cases of up to 40 percent.

This paper presented a strategy and validated it in an experiment with a Digital Twin. The next step is to implement this


strategy for the open-source scheduler Slurm [34] and verify the results on real hardware. This requires close coordination with the users of the HPC cluster as it affects the availability of the system. This paper relied on a node energy model with prior knowledge of the power profile of the applications. In a production scenario, the Digital Twin would need to learn new applications and use an estimate, e.g. of the very power intensive HPL benchmark, for unknown applications.

ACKNOWLEDGEMENT


We would like to thank Holger Duzinski and Thomas Várdu for their support of this research project. They provided us with helpful insight into the campus infrastructure and enabled access to sensors and meters without which the Digital Twin development would not be possible.

REFERENCES

- [1] Prometheus GmbH, “Top500 list,” May 2024. [Online]. Available: <https://top500.org/>
- [2] A. Konopelko, L. Kostecka-Tomaszewska, and K. Czerewacz-Filipowicz, “Rethinking eu countries’ energy security policy resulting from the ongoing energy crisis: Polish and german standpoints,” *Energies*, vol. 16, no. 13, 2023. doi: 10.3390/en16135132. [Online]. Available: <https://www.mdpi.com/1996-1073/16/13/5132>
- [3] Bundesministerium für Wirtschaft und Klimaschutz, “Ordinances on energy saving ensikumav and ensimimav,” Sep. 2022. [Online]. Available: <https://www.bmwk.de/Redaktion/DE/Downloads/E/ensikumav.html>
- [4] A. Kammeyer, F. Burger, D. Lübbert, and K. Wolter, “Towards an hpc cluster digital twin and scheduling framework for improved energy efficiency,” in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35, 2023. doi: 10.15439/2023F3797 p. 265–268.
- [5] P. Czarnul, J. Proficz, and A. Krzywaniak, “Energy-aware high-performance computing: Survey of state-of-the-art tools, techniques, and environments,” *Scientific Programming*, vol. 2019, p. 8348791, 2019. doi: 10.1155/2019/8348791. [Online]. Available: <https://doi.org/10.1155/2019/8348791>
- [6] B. Kocot, P. Czarnul, and J. Proficz, “Energy-aware scheduling for high-performance computing systems: A survey,” *Energies*, vol. 16, no. 2, 2023. doi: 10.3390/en16020890. [Online]. Available: <https://www.mdpi.com/1996-1073/16/2/890>
- [7] J. Corbalan and L. Brochard, “Ear: Energy management framework for supercomputers,” *Barcelona Supercomputing Center (BSC) Working paper*, 2019.



(a) Case of synthetic power limit function based on cyclic rising and falling limit. Absolute values for the two metrics mean job wait time and overall simulation time. (b) Case of real power limit function based on campus power usage. Absolute values for the two metrics mean job wait time and overall simulation time.



(c) Case of synthetic power limit function. Relative values for the two metrics mean job wait time and overall simulation time normalised to the first policy. (d) Case of real power limit function based on campus power usage. Relative values for the two metrics mean job wait time and overall simulation time normalised to the first policy.

Figure 6: Results of the experiments

- [8] M. D'Amico and J. C. Gonzalez, "Energy hardware and workload aware job scheduling towards interconnected hpc environments," *IEEE Transactions on Parallel and Distributed Systems*, p. 1, 2021. doi: 10.1109/TPDS.2021.3090334
- [9] B. Bylina, J. Bylina, and M. Piekarcz, "Impact of processor frequency scaling on performance and energy consumption for wz factorization on multicore architecture," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35, 2023. doi: 10.15439/2023F6213 p. 377–383.
- [10] B. Bylina and M. Piekarcz, "The scalability in terms of the time and the energy for several matrix factorizations on a multicore machine," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35, 2023. doi: 10.15439/2023F3506 p. 895–900.
- [11] A. Krzywaniak, J. Proficz, and P. Czarnul, "Analyzing energy/performance trade-offs with power capping for parallel applications on modern multi and many core processors," in *2018 Federated Conference on Computer Science and Information Systems (FedCSIS)*, 2018. doi: 10.15439/2018F177 p. 339–346.
- [12] H. van der Valk, H. Haße, F. Möller, and B. Otto, "Archetypes of digital twins," *Business & Information Systems Engineering*, vol. 64, no. 3, p. 375–391, Jun 2022. doi: 10.1007/s12599-021-00727-7. [Online]. Available: <https://doi.org/10.1007/s12599-021-00727-7>
- [13] ISO Central Secretary, "Digital twin – concepts and terminology," International Organization for Standardization, Geneva, CH, Standard ISO/IEC 30173:2023, Nov. 2023. [Online]. Available: <https://www.iso.org/standard/81442.html>
- [14] N. A. Simakov, M. D. Innus, M. D. Jones, R. L. DeLeon, J. P. White, S. M. Gallo, A. K. Patra, and T. R. Furlani, "A slurm simulator: Implementation and parametric analysis," in *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation*, S. Jarvis, S. Wright, and S. Hammond, Eds. Cham: Springer International Publishing, 2018. doi: 10.1007/978-3-319-72971-8_10. ISBN 978-3-319-72971-8 p. 197–217.

- [15] N. A. Simakov, R. L. Deleon, Y. Lin, P. S. Hoffmann, and W. R. Mathias, "Developing accurate slurm simulator," in *Practice and Experience in Advanced Research Computing*, ser. PEARC '22. New York, NY, USA: Association for Computing Machinery, 2022. doi: 10.1145/3491418.3535178. ISBN 9781450391610. [Online]. Available: <https://doi.org/10.1145/3491418.3535178>
- [16] A. Jokanovic, M. D'Amico, and J. Corbalan, "Evaluating slurm simulator with real-machine slurm and vice versa," in *2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, 2018. doi: 10.1109/PMBS.2018.8641556 p. 72–82.
- [17] T. Ohmura, Y. Shimomura, R. Egawa, and H. Takizawa, "Toward building a digital twin of job scheduling and power management on an hpc system," in *Job Scheduling Strategies for Parallel Processing*, D. Klusáček, C. Julita, and G. P. Rodrigo, Eds. Cham: Springer Nature Switzerland, 2023. doi: 10.1007/978-3-031-22698-4_3. ISBN 978-3-031-22698-4 p. 47–67.
- [18] J. M. Kunkel, H. Shoukourian, M. R. Heidari, and T. Wilde, "Interference of billing and scheduling strategies for energy and cost savings in modern data centers," *Sustainable Computing: Informatics and Systems*, vol. 23, p. 49–66, 2019. doi: 10.1016/j.suscom.2019.04.003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S221053791830297X>
- [19] X. Yang, Z. Zhou, S. Wallace, Z. Lan, W. Tang, S. Coghlan, and M. E. Papka, "Integrating dynamic pricing of electricity into energy aware scheduling for hpc systems," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC '13. New York, NY, USA: Association for Computing Machinery, 2013. doi: 10.1145/2503210.2503264. ISBN 9781450323789. [Online]. Available: <https://doi.org/10.1145/2503210.2503264>
- [20] A. Kammeyer, F. Burger, D. Lübbert, and K. Wolter, "Developing a digital twin to measure and optimise hpc efficiency." Submitted to IMEKO World Congress 2024, 2024.
- [21] Intel Corporation, "Intel xeon processor e5-2690 v4," Jul. 2024. [Online]. Available: <https://ark.intel.com/content/www/us/en/ark/products/91770/intel-xeon-processor-e5-2690-v4-35m-cache-2-60-ghz.html>
- [22] —, "Intel xeon gold 6132 processor," Jul. 2024. [Online]. Available: <https://ark.intel.com/content/www/us/en/ark/products/123541/intel-xeon-gold-6132-processor-19-25m-cache-2-60-ghz.html>
- [23] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, p. 114 ff., Apr. 1965. doi: 10.1109/NSSC.2006.4785860
- [24] D. G. Feitelson, D. Tsafir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, p. 2967–2982, 2014. doi: 10.1016/j.jpdc.2014.06.013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731514001154>
- [25] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary, "Hpl - a portable implementation of the high-performance linpack benchmark for distributed-memory computers," Dec. 2018, version 2.3. [Online]. Available: <https://www.netlib.org/benchmark/hpl/>
- [26] J. Dongarra, M. A. Heroux, and P. Luszczek, "High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems," *The International Journal of High Performance Computing Applications*, vol. 30, no. 1, p. 3–10, 2016. doi: 10.1177/1094342015593158. [Online]. Available: <https://doi.org/10.1177/1094342015593158>
- [27] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A tensorial approach to computational continuum mechanics using object-oriented techniques," *Computer in Physics*, vol. 12, no. 6, p. 620–631, 11 1998. doi: 10.1063/1.168744. [Online]. Available: <https://doi.org/10.1063/1.168744>
- [28] S. Agostinelli, J. Allison, K. Amako et al., "Geant4—a simulation toolkit," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, p. 250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>
- [29] N. Sudermann-Merx, *Fortgeschrittene Modellierungstechniken*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2023, p. 161–193. ISBN 978-3-662-67381-2. [Online]. Available: https://doi.org/10.1007/978-3-662-67381-2_7
- [30] D. Kolossa and G. Grübel, "Evolutionary computation and nonlinear programming in multi-model-robust control design," in *Real-World Applications of Evolutionary Computing*, S. Cagnoni, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000. ISBN 978-3-540-45561-5 p. 147–157.
- [31] T. Wilde, A. Auweter, and H. Shoukourian, "The 4 pillar framework for energy efficient hpc data centers," *Computer Science - Research and Development*, vol. 29, no. 3, p. 241–251, Aug 2014. doi: 10.1007/s00450-013-0244-6. [Online]. Available: <https://doi.org/10.1007/s00450-013-0244-6>
- [32] D. G. Feitelson, "Packing schemes for gang scheduling," in *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996. ISBN 978-3-540-70710-3 p. 89–110.
- [33] D. G. Feitelson and M. A. Jette, "Improved utilization and responsiveness with gang scheduling," in *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997. ISBN 978-3-540-69599-8 p. 238–261.
- [34] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple linux utility for resource management," in *Job Scheduling Strategies for Parallel Processing*, D. Feitelson, L. Rudolph, and U. Schwiegelshohn, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003. doi: 10.1007/10968987_3. ISBN 978-3-540-39727-4 p. 44–60.

Teaching Beginners to Program: should we start with block-based, text-based, or both notations?

Tomaž Kosar*, Srđa Bjeladinović†, Dragana Ostojčić*, Milica S. Škembarević†, Žiga Leber*, Olga A. Jejić†, Filip Furtula†, Miloš D. Ljubisavljević†, Ivan S. Luković†, Marjan Mernik*

* University of Maribor, Faculty of Electrical Engineering and Computer Science, Maribor, Slovenia
{tomaz.kosar, dragan.ostojic, ziga.leber, marjan.mernik}@um.si

† University of Belgrade, Faculty of Organizational Sciences, Belgrade, Serbia
{srdja.bjeladinovic, milica.skembarevic, olga.jejic, filip.furtula, milos.ljubisavljevic, ivan.lukovic}@fon.bg.ac.rs

Abstract—Teaching programming poses countless challenges. One of them is determining the most effective notation to introduce coding concepts to beginners. This paper examines the merits and drawbacks of introducing block-based, text-based, or both notations at the same time when it comes to learning basic programming concepts. By comparing these approaches, the objective of this research is to clarify and assess the learning outcomes related to teaching beginners through different notations. In this empirical study, we report on a controlled experiment during short-term visits that promoted programming in primary schools. Our multinational study divided participants into three groups, one using block-based, one using text-based, and one using both notations. After training, the participants were solving practical programming assignments. The study results revealed that the participants’ performance was not influenced by notation usage, as there was no statistical significance between the three groups. However, the performance outcomes were correlated with the duration of the sessions. Our findings from the controlled experiment suggest that educators can utilize different notations confidently while teaching beginners the first steps in programming.

I. INTRODUCTION

INTRODUCING programming to children in primary school is crucial in today’s digital world. Early exposure to programming not only equips children with valuable technical skills, but also enhances their complex problem-solving abilities [1], creativity, communication, and teamwork capabilities. Despite its importance, programming is still not a part of the primary school curriculum in many countries (e.g., Slovenia). This gap in primary school education can lead to an essential void in the future job markets. By integrating programming into the curriculum, we can ensure that all children have the opportunity to develop these essential skills, preparing them for a future where technology, as we all agree, will play an important role in different aspects.

To address this deficit in primary school education, we occasionally visit schools, spending time with children demonstrating software development. During these visits, we aim to engage children, by showcasing the development of simple

This work is sponsored by the bilateral project “Programming Environments with Simultaneous Multiple Representations in Support of Early Programming Education” between Slovenia and Serbia. The Slovenian authors acknowledge the financial support from the Slovenian Research Agency (Research Core Funding No. P2-0041). This research has been supported by the Faculty of Organizational Sciences of University of Belgrade.

applications and potential career opportunities in the field of Computer Science. We believe that, as guest lecturers in primary schools, we have several significant impacts. Guest lecturers can bring enthusiasm to the subject, which is different from the usual ones, making it more engaging for children. Moreover, guest lecturers can introduce diverse perspectives and knowledge that might not be readily available through the standard curriculum, enriching the children’s learning experience and highlighting the importance and relevance of programming skills in today’s digital world.

However, teaching programming to beginners is a complex task that requires careful consideration of various factors. One of the key challenges is the limited time available for teaching programming to children during visits. Another critical decision involves selecting the appropriate notation for introducing programming concepts. Using block-based notation (e.g., Scratch [2], App Inventor [3]), allows for the creation of visually appealing games and applications. Conversely, text-based notation enables the teaching of fundamental programming concepts. Each approach has its advantages and disadvantages.

On the other hand, the transition from block-based to text-based programming is often highlighted as problematic; starting with block-based notation can lead to novice programmers being reluctant to switch to textual notation. This is a common pitfall that programming educators encounter during short-term visits and in the regular curriculum.

To address this problem, educational tools have emerged that enable educators to teach novice programmers using both notations simultaneously [4], [5], [6]. The most significant advantage of this multi-representational environment is that the transition from block-based to text-based notation occurs very naturally. For example, consider how time-consuming it is to write math equations using blocks; in contrast, text-based notation allows for the expression of math equations in a more natural and efficient manner. This dual-notation strategy helps students integrate seamlessly and understand both forms of programming, easing the learning curve and enhancing their overall comprehension.

Although we developed one such multi-representational environment called Poliglot [6] that enables programming with both notations simultaneously, we are still determining if this duality impacts the mastery of basic programming concepts.

Therefore, this paper's motivation is to explore which notation, block-based (Scratch, App Inventor), text-based (e.g., Python), or dual approaches (e.g., Poliglot), enhances the understanding of basic programming concepts most effectively and establishes a solid foundation for children's interest in programming. More specifically, we were motivated by the following research question: Does specific programming notation affect the participants' test performance after training sessions during short-term visits to primary schools?

We designed a controlled experiment [7], [8], [9], conducted in classrooms during short-term visits (2 hours), including training and a brief test at the end of the sessions. This multinational study was executed in two different countries, providing a broad perspective on the effectiveness of these notations. Our multinational study divided the participants into three groups: one using block-based notation, another using text-based notation, and the third using both. After undergoing training, the participants were assigned practical programming tasks. The results of the study indicated that the type of notation used did not affect the participants' performance.

The paper is structured as follows: The second section provides crucial insights into the background of the three different notations/environments: block-based, text-based, and multi-representational environments (block-based and text-based notation presented simultaneously). The following section introduces the multi-representational environment Poliglot briefly, developed by one of the universities participating in this study. The fourth section reviews related studies. The fifth section illuminates the experimental design, goals, and data collection instruments. The sixth section presents the comparative results of the experimental study. The seventh section outlines and discusses the essential findings from the empirical investigation. Subsequently, the following section exposes critical threats to the validity of the results in this study. Finally, in the last section, conclusions are drawn regarding the research outcomes.

II. BACKGROUND

As stated earlier, our experiment included three notations. In this section we introduce these alternatives, and explain their dynamics and potential for their usage as an introduction to the programming world. Each of the presented tools was analyzed from the aspect of comprehensiveness to the novice programmers, i.e., children who had never encountered programming concepts before.

Block-based notation: Block programming allows early-age students to become familiar with the basic programming concepts while strengthening programming logic by applying visual components. Dedicated editors provide visualization of programming constructs and learning through play. Pure block notation is still often used for children who encounter programming at the earliest age. Block notation eliminates certain logical or semantic types of errors that occur in text editors, since blocks can be combined according to clearly defined rules in advance. By adopting the rules of combining blocks, the students learn to eliminate inevitable mistakes

spontaneously, contributing to faster acquisition of textual notation.

Scratch is an editor that allows learning block programming notation through play. It is based on a multi-panel, single-window setup, which provides transparency and clear visibility at all times. Every change is immediately noticeable, giving the impression that the program is "alive" [10]. Scratch supports hands-on, one of the main approaches to practicing coding [11], through the ease of testing each block and learning through changes and play. Although based on a block, Scratch can represent a reasonable basis for adopting a bottom-up approach in programming, but also for introducing students to extremely fine-grained programming [12].

Alice is another tool that enables active student engagement during the process of learning programming skills [13]. As the authors stated, one of the major challenges for novice students in programming is "putting the pieces together". Alice provides 3-D visualization for solving different programming problems. The animated programming environment in 3-D enables students to research further, and develop algorithms for animating objects in an even more intuitive and interesting way than 2-D environments provide. Creating methods for objects and testing them in the dynamic 3-D environment can enhance the adoption of object-oriented programming using textual notation.

To approach the young generations and activate them, not only when they are at the computer in the classrooms, but other systems can also be used, such as App Inventor. This tool is available on mobile devices. By using mobile devices, students can practice programming spontaneously and intuitively, even in moments of leisure. With this, through the game and constant availability, interest in programming can be accelerated less formally [12].

Text-based notation: Python is renowned as a multi-purpose programming language that can be utilized on various platforms [14]. The simple and minimalistic text-based syntax makes Python a convenient programming language for beginners, whereas various specialized modules that can be imported contribute to the versatility of this programming language. According to the data in [15], based on the number of Google searches for the tutorials, Python surpassed Java in 2018, and has been the most popular programming language ever since. Python enforces indentation as a way of separating nested blocks of code, which leads to a more visually intuitive way of reading and understanding the code (since indentation is considered a part of the syntax, and not just a recommendation in coding style, e.g., in Java). The role of parentheses is relatively reduced compared to Java or other object-oriented languages, where parentheses have the role of code separators and proprietorship indicators. Python is a high-level programming language with low-level machine instructions hidden from the developers, thus increasing comprehension and softening the learning curve. Another notable characteristic of Python is the dynamic assignment of variable type based on the given value of the variable, and there is no need for preemptive type declaration. Another

benefit of the Python programming language is the possibility of functionality extensions that can be achieved with the addition of predefined packages (modules) to the program. The separation of functionalities into modules contributes to the code's overall simplicity and reduces imports of unnecessary functionalities. In this way, the user does not need to be familiar with all functionalities at once, but can instead study module by module, depending on their needs.

Both notations - Multiple-Representation Environments:

Block-based languages are a popular way to introduce programming and create educational programming environments. However, users eventually need to transition to textual notation to develop more complex programs. Significant efforts have been made to aid this transition through various methods, including presenting translated versions, dual-mode, multiple-representation, and hybrid environments. Examples of these environments include such tools as Tiled Grace [4], Blockly [16], Pencil code [5], Droplet [17], Greenfoot [18], and Poliglot [6], all designed to address this challenge.

Let's introduce some of these environments briefly. Tiled Grace [4] is a tiled-based editor for the Grace programming language. Using tiles enables visualization of the code, and there is also support for text editing. Textual editing of code expands tiles' visualization, enabling seamless change of the work environment and easy transition between source code and visual representation of the code [4]. Droplet [17] is a library designed to create dual-mode environments. It translates code by inserting tags into the textual code, to indicate which parts will be represented as blocks. These blocks are then displayed by extracting the code between the tags and presenting it within the block structure. Tags are added using an external parser, and precedence is handled by a custom JavaScript function that inserts parentheses into the blocks. The transition back to text-based notation is done by removing the tags while keeping the parentheses intact. On the other hand, Blockly [16] is a multiple-representation, web-based environment with open access, targeting primarily novice programmers in Data Science. It uses Python for its text-based notation, facilitating a smooth transition from block-based to text-based programming for beginners.

III. POLIGLOT

An example of multiple-representation environments is also Poliglot¹ [6], developed by the Slovenian partners in this paper. Poliglot is an educational programming environment designed for beginners who are taking their first steps in programming. Our experiences teaching programming as guest lecturers in primary schools inspired the development. Previously, we often started with block-based languages like Scratch and App Inventor, which engage children in programming effectively. These tools allowed users to create functional games, mobile applications, and more quickly. However, when the capabilities of block-based languages were exhausted, transitioning to text-based languages became necessary.

¹<https://poliglot.um.si/>

This transition posed a significant challenge for novice programmers. They had to start with the basics again, taking much longer to reach the level of complexity they had achieved with block-based languages. We frequently observed that this shift led to a loss of enthusiasm among learners.

Poliglot addresses this issue by introducing both notations simultaneously from the outset. It helps beginners connect each block to its textual representation, easing the transition. By presenting both notations together, Poliglot blurs the boundaries between block-based and text-based programming.

As learners gain experience, they often find that expressing themselves in text-based notation becomes easier than using blocks. This transition happens naturally within the Poliglot environment.

An example of the Poliglot system is shown in Figure 1. In this example, users input two numbers. Children can choose to use either block-based or text-based notation. When working in block-based notation, the corresponding text-based code appears simultaneously in the tool's top-right corner. User input in the block-based environment results in real-time text-based code updates, and vice versa. As noted, arithmetic or logical equations often prompt beginners to switch to text-based notation naturally within the multiple-representation environment Poliglot [6].

Poliglot employs pretty-printing abstract syntax trees (AST), a standard task in language workbenches as described by Fowler in [19], and also utilized in MPS [20]. In these programming environments, the end-user is not editing the code directly, but rather the AST, which is the model underlying the code. Programs can be understood as trees—a hierarchy of constructs that form the language behind the code. Each editor in MPS is merely one projection of the same model, and a projectional editor can have multiple projections, or representations, of the same code. In this context, Poliglot offers two projections: a block-based editor and a text-based editor [6].

Note that we do not favor Poliglot as a multi-representation environment. Instead, we encourage other researchers to conduct similar experiments using comparable tools, such as Grace [4], Blockly [16], Pencil Code [5], Droplet [17], Greenfoot [18], etc. This will help to strengthen the results from this study.

IV. RELATED WORK

The authors in [21] performed a quasi-experimental study investigating how modality (block-based and text-based environment) impacts high school Computer Science students by conducting two classes at the same school through the same curriculum and the same teacher using either the block-based or text-based programming environment (The Pencil.cc environment was used, which supports both modalities, but students were able to use only one modality). The outcome of this study [21] shows that the students' conceptual knowledge had been improved in both groups. However, the students using a block-based environment showed significant learning gains, as well as a higher attitude toward future programming




Fig. 1. Multiple-representation environment Poliglot with Limpid language

courses. On the other hand, no difference was found in both groups with respect to confidence and enjoyment. This work was later extended in [22], where the authors checked a hypothesis that gains in attitudinal and conceptual learning using a block-based environment would transfer to a conventional text-based programming language (Java). The study showed that, whilst students had a greater conceptual learning gain using a block-based environment, this was not transferred to the environment using the professional Java programming language. Furthermore, no difference in programming practices or attitudinal shifts was found between both groups. As such, this study [22], is important to show the limitations of block-based programming.

The study [23] tried to answer the difference between block-based and text-based environments on novice Com-

puter Science students' cognitive (knowledge, comprehension, application, analysis, synthesis, evaluation) and attitudinal (satisfaction, confidence, motivation, appreciation, enthusiasm) outcomes by performing a meta-analysis, which showed that block-based environments had a small effect on cognitive outcomes, and only a trivial effect on attitudinal outcomes.

The study [21] was extended in [24] by a third group using a hybrid block/text environment, with the main goal of how modality (block-based, text-based, hybrid block/text) influences programming practices (e.g., the number of runs, patterns in novice's help-seeking behavior). While the authors didn't find hybrid block/text modality superior, they did find some new programming practices. However, cognitive and attitudinal outcomes have not been measured and discussed. Our study extends this one [21], and brings additional evidence

in this field.

Students' difficulties in the transition from block-based to text-based environment have been discussed in [25], where it was shown that the students struggled to solve a new coding challenge in a text-based environment due to difficulties of one or more following aspects: readability, memorization of commands, memorization of syntax, native language of programming, typing/spelling and writing expressions.

The authors in [26] presented a Systematic Literature Review (SLR) [27] on the characteristics of block-based environments, and how block-based environments support learners in the transition to text-based programming, where the following distinct approaches were identified: Blocks-only, One-way Transition, Dual-modality, and Hybrid. Among those, dual-modality programming environments are the most effective for supporting students' transitioning to text-based programming.

V. EXPERIMENT DESIGN

The primary objective of this study is to examine the outcomes of assessments when they engage with programming using three distinct types of programming environments: block-based, text-based, and a combination of both notations. By comparing the results from these approaches, the study seeks to elucidate the learning outcomes associated with teaching novices using different programming notations.

The central hypothesis is that the type of notation used affects the participants' performance significantly. This hypothesis is based on the widespread use of block-based notation among educators teaching children the first steps in programming.

A. Participants

The experiment was multinational and multi-institutional, involving participants from two different countries: 236 from Serbia and 64 from Slovenia. The participants were primary school children aged between 11 and 14 years. The participants were sixth, seventh, and eighth-grade elementary school students. There were 84 sixth-grade students, 96 seventh-grade students, and 56 eighth-grade students from Serbia. In Slovenia, the participants were also elementary school students from the sixth to eighth grade. There were 30 sixth-grade, 29 seventh-grade, and four eighth-grade students. No prior participant selection was conducted for this study, resulting in the inclusion of participants with diverse backgrounds, varying levels of knowledge and experience in programming, and differing levels of interest in the subject. The assessment of previous knowledge and experience was based on the grades participants had at the end of a previous school year. The average grade for mathematics was 3.8 out of 5, and 4.9 out of 5 for informatics but only 162 participants had that subject.

B. Procedure

Each execution consisted of a background questionnaire, a programming class, and a final programming test, all within a two-hour timeframe. The ideal duration of the experiment would be four hours, during which the candidates would have

more time to try and experiment with the tool themselves, but the duration of the experiment was limited because of practical reasons; according to the schools' schedules Informatics classes last for 1 hour, and even merging two classes was causing inconvenience to the teachers. The background questionnaire assessed the participants' prior programming experience. Before the lectures, the participants were all given the same entrance questionnaire. Apart from the elementary questions, such as grade and gender, additional data about participants were gathered through various categories:

- 1) Previous experience with programming (formal and informal);
- 2) Computer interaction frequency;
- 3) Inclination to problem-solving tasks;
- 4) Mathematical knowledge, and
- 5) Level of logical thinking applied to solving problems.

Following the survey, the test group attended a lecture featuring a presentation that provided training through a specific tool and notation. The lecture focused on fundamental programming concepts, such as statements (e.g., printing), logic, arithmetic equations, variables, conditional statements, and loops. The participants followed primarily the educator's actions displayed on a large screen. After learning the basic programming concepts, the participants completed a final test comprising six questions that covered the topics taught during the training. All the questions used the same programming notation as the training session. Even though the questionnaires were anonymous, the results collected at the end of the lectures were paired with the data collected from the entry questionnaire, increasing the research coverage and exploring the inclination to programming concepts and way of thinking.

Multiple iterations of the study were conducted in both countries. For each new iteration, a different programming notation was used: block-based, text-based, and, finally, a combination of both notations. During the training sessions, all the educators utilized the same PowerPoint slides, which included explanations of the concepts, tasks, and correct program examples. This approach ensured consistency in the training provided by different educators. Prior to the main experiment, pilot studies were conducted to refine the background questionnaire, training materials, and the final test. Each question on the test offered five potential answers, with only one being correct.

C. Data Collection Instruments

The tests that were handed out after the lectures consisted of the same set of questions written in the corresponding notation based on the materials that were presented during the lectures (block-based Poliglot, block-based Poliglot combined with text notation, or text-based Python). The questions appeared in ascending order according to their difficulty. There were five question types:

- 1) Prediction of the given code execution: The participants were presented with a few lines of code with options on what the result of the execution of that code would be.

- 2) Finding a redundant piece of code: Based on the given code block, participants were asked to identify a redundant line of code that did not influence the program's execution.
- 3) Code insertion: A block of code was presented with one line missing; the participants were expected to select the line that would complete the block of code and provide a logical solution to the problem.
- 4) Identification of the logical errors in the code: The block of code was shown with a notice stating that there was a logical error in the code that needed to be identified.
- 5) Code modification: The last task required participants to change a line of code, thus changing the result of the code execution to match the description of the desired code behavior.

Even though the tests were written in different notations, the logic behind the question remained the same without any changes to the formulation of the question except the syntax. Figure 2 shows a question from a block-based test given to the participants after training. In this question, the participants were asked about the result of running the block-based program. This question is an example of the "Prediction of the given code execution" question type.

6

What is the result (output) of the program below? *

Program

Logic

Console

Variables

Numbers

program

a 5

b 4

c 2

if a < b

print a

else if b < a

print b

else

print c

a) 2

b) 4

c) 5

d) 542

e) Nothing

Fig. 2. Question from block-based test

Half of the questions on the test for multiple-representation environments were written in the text notation, and the other half of the questions were written using the block notation. The tests were prepared in two languages, Slovenian and Serbian, allowing the participants to solve tasks in their native language. The points awarded to each question were also the same (each answer was worth 1 point) to make the tests comparable.

VI. RESULTS

This section compares the participants' performance from two different countries (Serbia is referred to as country 1

and Slovenia is referred to as country 2) with three different notations.

A. Comparative Comparison: between-subjects study

Block-based vs. Text-based results (between-subjects design) - country 1

The first comparison examined the test results of the Serbian participants (country 1) who attended lectures and took the tests in block notation (Poliglot), versus those who did so in text notation (Python). After data cleansing, there were 52 valid responses for the block-based Poliglot test and 65 valid responses for the text-based Python test. Some tests could not be paired with their initial counterparts, and were therefore discarded. Additionally, instances of double submissions by the same participant reduced the number of tests considered for analysis further.

Statistical testing was conducted on these data, with a significance threshold set at $\alpha = .05$. The Shapiro-Wilk test was used to check for normal distribution. Since the data deviated from a normal distribution, the non-parametric Mann-Whitney U test was employed to compare the two independent samples. The slight difference in the mean scores (see Table I) between the two groups was not statistically significant (p-value = 0.621).

Block-based vs. text-based results: country 2

To verify the consistency of the results obtained in Serbia, another between-subjects study was conducted in Slovenia (country 2). Table II presents the performance results, measured as the percentage of correct responses, to assess programming knowledge after training. Both Group I (block-based) and Group II (text-based) completed an equal number of tasks with identical question types and complexity. An examination of Table II reveals that the text-based group outperformed the block-based group, as indicated by the mean scores (36,67% vs. 34,62%).

Once again, the data deviated from a normal distribution, necessitating the non-parametric Mann-Whitney U test to compare the two independent samples. Despite the observed difference in mean scores between the two groups, this difference was not statistically significant (p-value = 0.831).

These results are consistent with our findings from the initial study conducted in Serbia. Similarly, they aligned with those obtained in the study by Weintrop et al. [22], which demonstrated that, when participants start programming in either block or text notation, there is no difference in correctness or efficiency when they transition to a professional text-based language. This study underscores that the initial programming notation does not impact subsequent performance.

Multiple-Representation Environments vs. Text-based results (between-subjects design) - country 1

The second comparison in this research examined a multiple-representation environment (Poliglot, featuring both text and block notation) against a text-based environment (Python). A total of 39 participants took the test using Poliglot

TABLE I
PERFORMANCE RESULTS: BLOCK-BASED VS. TEXT-BASED (MANN-WHITNEY U TEST) – COUNTRY 1

Part	Mean	N	Std. Dev.	Median	Mean Rank	Z	p-value
Group I (block-based)	45,67	52	22,40	50,00	57,29	-0,495	0,621
Group II (text-based)	48,08	65	23,72	50,0	60,37		

TABLE II
PERFORMANCE RESULTS: BLOCK-BASED VS. TEXT-BASED (MANN-WHITNEY U TEST) – COUNTRY 2

Part	Mean	N	Std. Dev.	Median	Mean Rank	Z	p-value
Group I (block-based)	34,62	26	23,53	33,33	25,58	-0,213	0,831
Group II (text-based)	36,67	25	20,41	33,33	95,59		

(see Table III). Among them, some participants did not answer any questions correctly, and none achieved the maximum score. The data for the text-based environment (a total of 65 participants) were the same as shown in Table I, but were then compared statistically against the data collected from the participants who experienced both (block and text) notations side by side during their training.

The average score on the combined test was the highest of all three groups, at 49.04%. The Standard Deviation for this test was the same as for the text-based test, with a value of 23.71%. The variance also matched that of the text-based test, which was 23.72%. Once again, the difference between the groups was not statistically significant (p -value = 0.775), as determined by the Mann-Whitney U statistical test.

B. Improvement Comparison: within-subjects study

Another interesting perspective was examining how studying both notations simultaneously affected performance in a text-based notation. It is important to note that the text-based notations differed – in Poliglot, we used our text language, Limpid, while, in the text-based experiment, we used Python. This study refers to a within-subjects design, meaning that each participant first worked in a multiple-representation environment (Poliglot) and then with the text-based notation (Python).

Additionally, due to the poor initial results, we extended our study to include a 4-hour training session for each treatment. This extension provided more time for repetition, and allowed the participants ample time to develop their solutions independently without time pressure.

Table IV presents the performance results. Both Group I (using both notations) and Group II (using text-based notation) outperformed all the previous executions (e.g., see Mean Column in Table II). A closer examination revealed that the participants' performance in Python notation gave better results than Poliglot, with mean scores of 60.0% compared to 55.55%, respectively. However, the difference was not statistically significant (p -value = 0.813).

VII. DISCUSSION

To understand the outcomes of our controlled experiment better, we present the results of a study utilizing a background questionnaire, focusing particularly on the Informatics and Mathematical backgrounds of the participants involved.

Two groups of between-subject tests were conducted in this paper. The first one was block-based vs. text-based. We cannot confirm the differences in the results statistically. However, a higher variance of deviations indicates a greater dispersion of achieved points for textual notation. The fact that each participant achieved at least 1 point also speaks in favor of the block notation, which was not the case with the text notation, because there were two participants with 0 points in the sample.

In favor of the uniformity of the useful prior knowledge of the participants in the two tested groups, the almost insignificant differences in the average grade in Informatics that the students of both tested groups had (4.98 students who did block notation, 4.93 students who did textual notation), as well as the successfulness of resolving three logical tasks given in the input survey (the students who did the block notation averaged 1.96 points from 3, while the students who did the text notation had 1.97 on average). A slightly higher average grade in Mathematics was present in the students who did the block notation (3.87) compared to the students who did only the textual notation (3.67). However, the average score was on the side of textual notation, so it can be concluded that prior knowledge of Mathematics and Informatics was not crucial for the achieved result in programming in the tested groups.

The second group of between-subject tests was multiple-notation vs. text-based. Identical values of Standard Deviation and variance in both approaches suggest the uniformity of students within the sample, although the number of students with combined notation was the smallest of all three tested groups. Furthermore, the obtained results suggested a slight advantage of the combined notation compared to the textual one, which was represented by a slightly higher average grade, but also by the values of the first and third quartiles, while the fourth belonged to the textual notation, because there was no maximum number of points in the combined one. Based on the sample, it can be concluded that the students using block notation achieved better performance more easily, with lower and medium performance on the test. In contrast, for maximum performance, textual notation still had an edge. All of the above leads to the conclusion that Poliglot achieves an advantage by using both notations, and improves the average performance compared to exclusively block notation. Similar to the previously analyzed group of tests (block notation vs

TABLE III
PERFORMANCE RESULTS: MULTIPLE-REPRESENTATION ENVIRONMENTS VS. TEXT-BASED (MANN-WHITNEY U TEST) – COUNTRY 1

Part	Mean	N	Std. Dev.	Median	Mean Rank	Z	p-value
Group I (both notations)	49,04	39	23,71	50,00	53,83	-0,286	0,775
Group II (text-based)	48,08	65	23,72	50,00	51,85		

TABLE IV
PERFORMANCE RESULTS: MULTIPLE-REPRESENTATION ENVIRONMENTS VS. TEXT-BASED (MANN-WHITNEY U TEST) – COUNTRY 2

Part	Mean	N	Std. Dev.	Median	Mean Rank	Z	p-value
Group I (both notations)	55,55	12	22,85	50,00	11,21	-0,237	0,813
Group II (text-based)	60,00	10	31,62	58,34	11,85		

textual notation), when comparing the groups of students who did combined notation vs textual notation, it can be concluded that the average marks in Informatics (an average of 4.87 for the students with combined notation vs 4.94 for the students with textual notation) and Mathematics (3.64 was the average grade for the students learning combined notation vs 3.67 for the students learning textual notation) were quite uniform. A slight difference in the average number of points on the logical tasks of the input survey was evident in favor of the students who did text notation (average 1.97 from 3 points vs average 1.77 from 3 points). Despite this, the students with combined notation achieved the best results of average points on the output of all three analyzed groups in Serbia, which can lead to the conclusion that, by applying both notations, there is scope for achieving better performance, even for those students who are closer to block notation (e.g. who have experience in using Scratch or some other environments), but also among students who are closer to textual notation, precisely because of the possibility of choosing a notation that is more convenient for them.

VIII. THREATS TO VALIDITY

This section discusses the construct, internal, and external validity threats [28] associated with our experiment.

A. Construct Validity

In our experiment, we aimed to measure the effect of notation on test results. The participants were assessed with multiple-choice questions after two hours of training in specific notations: block-based, textual, and both notations simultaneously. The use of multiple-choice questions may have influenced the results. Different outcomes might have emerged if we had used code implementation or code completion questions instead.

Another potential threat to construct validity was the complexity of the questions. The test results were generally low, with almost all the experiments resulting in an average performance of 50% or less. The outcomes of our experiments might differ if the question complexity were reduced.

The training sessions were limited to two hours, as requested by the primary schools. This constraint might have influenced our results significantly. To investigate this, we conducted one session at the university, during which the training duration

was extended to four hours. This extended training included additional functionality and repetitions of mastering the same programming concept. Consequently, we observed higher performance results (see Table IV). However, these improvements were seen in both the multiple-representation and text-based groups, and the outcomes were not statistically significant.

We used Poliglot for block-based and multiple-representation environments, although alternative tools exist for both notations (e.g., BlockPy). Our experiments did not include the usage of these other tools, so our findings are specific to Poliglot.

B. Internal Validity

One potential threat to internal validity is the quality of instruction provided to novice programmers during our experiment. Although we standardized the training materials (such as presentations), the use of different lecturers may have influenced the outcomes of our experiment.

The sample size may have influenced the results of the within-subjects study, as only one execution was conducted, with 10 participants completing both tests. To enhance the reliability and validity of these findings, it is crucial to perform multiple repetitions of the experiment with a larger sample size. This will help mitigate potential biases, and provide a more robust understanding of the observed effects. Despite the small sample size, the extended training period within that study clearly demonstrated a positive influence on the results.

Another concern for internal validity is the possibility of cheating during the tests, which could compromise the results. This is a common issue in educational settings, particularly when tests are administered in classroom environments.

C. External Validity

Our experiment's specific context and settings might influence our study's external validity. The results could vary with different participant demographics, educational environments, or levels of prior programming experience among the participants. Our findings were derived from a small set of schools in two countries. To generalize these findings, further research is needed, involving more institutions and conducting multi-institutional and multinational studies in diverse settings.

IX. CONCLUSION

By evaluating notation's impact on learning outcomes, this paper aims to provide insights into using three distinct approaches for teaching programming to novices.

In conclusion, our study examines the choice of notation, whether text-based, block-based, or multi-representational environments, for teaching programming concepts to novice programmers during short-term visits to primary schools. Our findings demonstrate that the choice of notation (block-based, text-based, or both) did not result in significant deviations in the participants' outcomes.

For future work, longitudinal studies and follow-up research are essential, to explore the potential effects of teaching various notations in greater depth. We are planning additional experiment repetitions [29] with the same and similar settings to validate our current findings, ensuring greater accuracy and reliability. In this study, we demonstrated how performance outcomes correlated with the duration of the training (2 hours vs. 4 hours). The sample size in the 4 hour- training session was small (see Table IV, again). Therefore, we intend to test our findings with future experiments. With a larger sample size, we can also analyze the participants' results with similar backgrounds, the same age, and the same conditions, thereby isolating the variables to be evaluated. Extending a multinational approach and considering diverse experiment settings is essential for a comprehensive understanding of using different programming notations and environments for teaching novice programmers.

REFERENCES

- [1] B. Bubnič, M. Mernik, and T. Kosar, "Exploring the predictive potential of complex problem-solving in computing education: A case study in the introductory programming course," *Mathematics*, vol. 12, no. 11, 2024. [Online]. Available: <https://www.mdpi.com/2227-7390/12/11/1655>
- [2] M. Resnick, J. Maloney, A. Monroy-Hernández, N. Rusk, E. Eastmond, K. Brennan, A. Millner, E. Rosenbaum, J. Silver, B. Silverman *et al.*, "Scratch: programming for all," *Communications of the ACM*, vol. 52, no. 11, pp. 60–67, 2009.
- [3] D. Wolber, H. Abelson, E. Spertus, and L. Looney, *App inventor*. O'Reilly Media, Inc., 2011.
- [4] M. Homer and J. Noble, "A tile-based editor for a textual programming language," in *IEEE Working Conference on Software Visualisation (VISSOFT)*, 2013, pp. 1–4.
- [5] D. Bau, D. A. Bau, M. Dawson, and C. S. Pickens, "Pencil code: block code for a text world," in *Proceedings of the 14th International Conference on Interaction Design and Children*, 2015, pp. 445–448.
- [6] Ž. Leber, M. Črepinek, and T. Kosar, "Simultaneous multiple representation editing environment for primary school education," in *2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 2019, pp. 175–179.
- [7] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in software engineering*. Springer Science & Business Media, 2012.
- [8] T. Kosar, M. Mernik, and J. C. Carver, "Program comprehension of domain-specific and general-purpose languages: comparison using a family of experiments," *Empirical software engineering*, vol. 17, pp. 276–304, 2012.
- [9] L. Alves, D. Gajić, P. Rangel Henriques, V. Ivančević, V. Ivković, M. Lalić, I. Luković, M. J. Varanda Pereira, S. Popov, and P. Correia Tavares, "C tutor usage in relation to student achievement and progress: A study of introductory programming courses in Portugal and Serbia," *Computer Applications in Engineering Education*, vol. 28, no. 5, pp. 1058–1071, 2020.
- [10] J. Maloney, M. Resnick, N. Rusk, B. Silverman, and E. Eastmond, "The Scratch programming language and environment," *ACM Transactions on Computing Education*, vol. 10, no. 4, pp. 138–144, 2010.
- [11] V. Handur, P. D. Kalwad, M. S. Patil, V. G. Garagad, P. Yeligar, Nagaratna and Pattar, D. Mehta, P. Baligar, and J. H., "Integrating class and laboratory with hands-on programming: Its benefits and challenges," in *IEEE 4th International Conference on MOOCs, Innovation and Technology in Education (MITE)*, 2016, pp. 163–168.
- [12] O. M. Salant, M. Armoni, and M. Ben-Ari, "Habits of programming in Scratch," in *ITiCSE '11: Proceedings of the 16th annual joint conference on Innovation and technology in computer science education*, 2011, pp. 168–172.
- [13] S. Cooper, W. Dann, and R. Pausch, "Alice: a 3-d tool for introductory programming concepts," *Journal of Computing Sciences in Colleges*, vol. 15, no. 5, pp. 107–116, 2000.
- [14] A. Martelli, A. M. Ravenscroft, S. Holden, and P. McGuire, *Python in a Nutshell*. " O'Reilly Media, Inc.", 2023.
- [15] PYPL, "PYPL - popularity of programming language," <https://pypl.github.io/PYPL.html>, accessed: 22.05.2024.
- [16] A. C. Bart, J. Tibau, E. Tilevich, C. A. Shaffer, and D. Kafura, "BlockPy: An open access data-science environment for introductory programmers," *Computer*, vol. 50, no. 5, pp. 18–26, 2017.
- [17] D. Bau, "Droplet, a blocks-based editor for text code," *Journal of Computing Sciences in Colleges*, vol. 30, no. 6, pp. 138–144, 2015.
- [18] M. Kölling, "The Greenfoot Programming Environment," *ACM Transactions on Computing Education (TOCE)*, vol. 10, no. 4, pp. 1–21, 2010.
- [19] M. Fowler, "Language Workbenches: The Killer-App for Domain Specific Languages? <http://www.martinfowler.com>," 2005.
- [20] M. Voelter and V. Pech, "Language modularity with the MPS language workbench," in *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 2012, pp. 1449–1450.
- [21] D. Weintrop and U. Wilensky, "Comparing block-based and text-based programming in high school computer science classrooms," *ACM Trans. Comput. Educ.*, vol. 18, no. 1, oct 2017. [Online]. Available: <https://doi.org/10.1145/3089799>
- [22] D. Weintrop and U. Wilensky, "Transitioning from introductory block-based and text-based environments to professional programming languages in high school computer science classrooms," *Computers & Education*, vol. 142, p. 103646, 2019.
- [23] F. T. Zhen Xu, Albert D. Ritzhaupt and K. Umamathy, "Block-based versus text-based programming environments on novice student learning outcomes: a meta-analysis study," *Computer Science Education*, vol. 29, no. 2-3, pp. 177–204, 2019. [Online]. Available: <https://doi.org/10.1080/08993408.2019.1565233>
- [24] D. Weintrop and U. Wilensky, "How block-based, text-based, and hybrid block/text modalities shape novice programming practices," *International Journal of Child-Computer Interaction*, vol. 17, pp. 83–92, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2212868917300314>
- [25] C. V. Alejandro Espinal and V. Guerrero-Bequis, "Student ability and difficulties with transfer from a block-based programming language into other programming languages: a case study in Colombia," *Computer Science Education*, vol. 33, no. 4, pp. 567–599, 2023. [Online]. Available: <https://doi.org/10.1080/08993408.2022.2079867>
- [26] Y. Lin and D. Weintrop, "The landscape of block-based programming: Characteristics of block-based environments and how they support the transition to text-based programming," *Journal of Computer Languages*, vol. 67, p. 101075, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S259011842100054X>
- [27] T. Kosar, S. Bohra, and M. Mernik, "A Systematic Mapping Study driven by the margin of error," *Journal of Systems and Software*, vol. 144, pp. 439–449, 2018.
- [28] R. Feldt and A. Magazinius, "Validity threats in empirical software engineering research - an initial survey," in *22nd International Conference on Software Engineering & Knowledge Engineering (SEKE'2010), Redwood City, San Francisco Bay, CA, USA, July 1 - July 3, 2010*. Knowledge Systems Institute Graduate School, 2010, pp. 374–379.
- [29] T. Kosar, S. Gaberc, J. C. Carver, and M. Mernik, "Program comprehension of domain-specific and general-purpose languages: replication of a family of experiments using integrated development environments," *Empirical Software Engineering*, vol. 23, pp. 2734–2763, 2018.

Pareto Optimal Solutions of the Biobjective Minimum Length Minimum Risk Spanning Trees Problem

Lasko M. Laskov

0000-0003-1833-818

Informatics Department

New Bulgarian University

21 Montevideo Str., 1618 Sofia, Bulgaria

Email: llaskov@nbu.bg

Marin L. Marinov

0009-0003-9544-819X

Informatics Department

New Bulgarian University

21 Montevideo Str., 1618 Sofia, Bulgaria

Email: mlmarinov@nbu.bg

Abstract—We propose an exact method that finds the complete Pareto front of the biobjective minimum length minimum risk spanning trees problem. The proposed method consists of the solution of two problems. The first problem is to compute a list of all minimum spanning trees with respect of the length criterion. The second problem is to construct the complete Pareto front itself, based on the list of all minimum spanning trees, found using the solution of the first problem.

We prove mathematically the correctness of all proposed algorithms, and we discuss their computational complexity. We also illustrate the presented solution with detailed numerical examples.

I. INTRODUCTION

THE minimum spanning tree problem is a classical problem in combinatorial optimization and graph theory. It is considered important because its numerous applications in telecommunication, electronics, clustering algorithms and pattern recognition. The classical well-known algorithms that solve this problem are proposed by Kruskal [1], Prim [2] and Dijkstra [3]. Since then, many algorithms have been published that solve the fundamental version of the problem [4] and its generalization [5].

Even though the solution of the single-criterion problem is essential for the field, many problems in practice require the construction of minimum spanning trees in networks with multiple objective functions, and in particular networks with two objective functions. Both exact and heuristic methods are developed to address the problem of biobjective minimum spanning trees.

In [6] authors propose a method that finds all Pareto optimal solutions of a biobjective minimum spanning trees problem that is based on computation of extreme efficient solutions and then the computation of the remaining non-extreme solutions. The second phase of the method is based on k-best algorithm that is sped up using heuristic enhancement.

In [7] is proposed a multi-objective metaheuristic approach to solve biobjective spanning trees with minimum total cost and minimum diameter problem. The described method is based on multi-objective evolutionary algorithm and on a

nondominated sorting genetic algorithm. Solution of the same version of the problem is proposed by [8], but this time the proposed method is exact, with both correctness and running time verified experimentally.

In this paper we propose an exact method that finds the complete Pareto optimal front of the biobjective minimum length minimum risk spanning trees problem. Our solution is based on the solution of two problems: (i) construct the list of all minimum spanning trees with respect of the length criterion; (ii) construct the complete list of all Pareto optimal solutions. Our algorithm that solves the first problem is an extension of the Prim's algorithm [2]. It finds a single minimum spanning tree according to the length criterion, with simultaneously completing a list of subproblems that define all the remaining minimum spanning trees in the network. The solution of the second problem is based on an algorithm that composes a list of all Pareto optimal trees that have minimum length.

Even though the classical minimum spanning tree problem is solved in polynomial computational complexity time (see for example [9]), the problem of calculation of the complete Pareto optimal front of the biobjective minimum spanning trees has exponential complexity and is classified as an NP-hard problem [6]. The exact solution that we propose in this paper depends on the number of classes of Pareto optimal solutions and on the number of minimum spanning trees with respect to the length criterion.

This paper is organized as follows. In Sec. II we provide basic notations and definitions used in the paper. In Sec. III we describe the algorithm that computes the list of all minimum spanning trees with respect to the length criterion. In Sec. IV we give the solution of the main problem of this paper and we describe the algorithm that constructs the complete Pareto front. Finally, in Sec. V we provide our conclusions.

II. BASIC NOTATIONS AND DEFINITIONS

Let $G = (V, E)$ is a connected undirected graph with $n = |V|$ number of vertices, and $m = |E|$ number of

edges. Without loss of generality we will assume that $V = \{1, 2, \dots, n\}$ and $E \subseteq V^2$.

We define the following two objective functions: $f : E \rightarrow \mathbb{R}_+$ and $g : E \rightarrow \mathbb{R}_+$.

The function $f : E \rightarrow \mathbb{R}_+$ assigns to each edge $(u, v) \in E$ the positive real number $f(u, v)$, which we will call *length* of the edge $e = (u, v)$. The function $g : E \rightarrow \mathbb{R}_+$ assigns to each edge $(u, v) \in E$ the positive real number $g(u, v)$, which we will call *risk* of the edge $e = (u, v)$.

With N we denote the network $N = (V, E, f, g)$. We assume that the network N is given using the adjacency lists Adj (see Eq. (1)). We will treat the adjacency lists of the network as its attribute and we adopt the notation $N.Adj$.

$$N.Adj = [[(v, f(u, v), g(u, v)), \dots], \dots]. \quad (1)$$

The adjacency lists are represented as an array of lists, where each element of the array $N.Adj[u]$ corresponds to a vertex $u \in V$ of the network and it stores a list composed of ordered triples of the type $(v, f(u, v), g(u, v))$, where $v \in V$ is an adjacent vertex, $f(u, v)$ is the length and $g(u, v)$ is the risk of the edge (u, v) .

For each subset of edges $A \subseteq E$, where

$$A = \{(u_1, v_1), \dots, (u_k, v_k)\}, \quad (2)$$

we define the following two numbers $x(A)$ and $y(A)$.

$$x(A) = \sum_{i=1}^k f(u_i, v_i) \quad (3)$$

$$y(A) = \max\{g(u_i, v_i) : i \in \{1, 2, \dots, k\}\} \quad (4)$$

We call the number $x(A)$ *length* of the subset of edges A , and we call the number $y(A)$ *risk* of the subset of edges A .

Also, if the set A given in (2) is acyclic and if it contains all the vertices of the network N , we will call it a *spanning tree* of N .

With W we denote the set of all spanning trees of the network N .

Definition 1. We call the tree $t' \in W$ Pareto optimal when there does not exist another tree $t \in W$, for which any of the following two conditions is fulfilled:

- $x(t) < x(t')$ and $y(t) \leq y(t')$;
- $x(t) \leq x(t')$ and $y(t) < y(t')$.

Definition 2. We will say that the tree t' and t are equivalent and we will denote it with $t' \sim t$, when $x(t') = x(t)$ and $y(t') = y(t)$.

Definition 3. We will say that the tree t is dominated by the tree t' denoted $t \prec t'$, when $x(t') < x(t)$ and $y(t') \leq y(t)$ or $x(t') \leq x(t)$ and $y(t') < y(t)$.

We denote with P the set of Pareto optimal trees in the network N . It is clear that

$$P = \bigcup_{i=1}^K P_i, \quad (5)$$

where P_i are the classes of Pareto equivalent optimal trees, and K is the number of such classes.

III. LIST OF ALL MINIMUM LENGTH SPANNING TREES

In this section we present the method to construct all minimum spanning trees with respect of the length criterion, without actually constructing other spanning trees. The main part of the method is an inductive procedure that traverses a single minimum spanning tree A , starting from its root (Procedure 2). Each iteration of the procedure discovers all the possibilities by which the so far traversed part of A can be completed to another spanning tree of minimum length. These possibilities are called *subproblems* below. The procedure stops when the entire minimum spanning tree A is constructed. In other words, we solve the following Problem 1.

Problem 1. Given the network $N = (V, E, f, g)$, find one minimum spanning tree A and compose a list of the subproblems that define the remaining minimum spanning trees in the network N .

We give the solution of the above Problem 1 with the Algorithm 3 which is an extension of the classical Prim's algorithm [2] that finds a single minimum spanning tree in a connected, undirected, weighted graph. The proposed Algorithm 3 simultaneously constructs the minimum spanning tree and composes the list of the subproblems that define the remaining minimum spanning trees in the network N . This strategy is possible because for the minimum spanning tree problem a strong version of the greedy principle, known as *greedy-choice property* [9], is fulfilled, or in other words, the Theorem 2 holds.

In analogy to the Prim's algorithm, using an inductive procedure we discover the edges of a minimum spanning tree and we include them consecutively into a set A . Initially, the set A is the empty set. In Algorithm 3, instead of the set A itself, we explicitly maintain only its storage using the vector *tree*. When $A = \emptyset$, then *tree* is the zero vector with n components. The inclusion of an edge (u, v) into A we store in *tree* by $tree[v] = u$. In this case we say that for the currently visited vertex v the vertex u is its parent. We determine the edges that we will include into the set A by traversing the set of vertices V . With Q we denote the set of vertices which are not yet visited by the algorithm. In the initial state, $Q = V$. With U we denote the set of vertices that are already visited. Apparently, $U = V \setminus Q$.

We choose an arbitrary vertex r for the root vertex of the spanning trees. Without loss of generality, we can select $r = 1$. With the selected root r we define $U = \{r\}$ and $Q = V \setminus U$. Then, one minimum spanning tree can be discovered using the following Procedure 1.

Procedure 1. Inductive procedure to construct single minimum spanning tree.

I. Base step.

- 1) Discover a vertex v_1 from the set Q for which

$$f(r, v_1) = \min\{f(r, v) : v \in Q\}. \quad (6)$$

- 2) Move the vertex v_1 from Q into the set U resulting in the new U and Q after the visiting of the vertex v_1 .

II. Inductive step.

- 1) Discover the vertices v_1 and v_2 , for which the conditions (7) and (8) are fulfilled.

$$v_1 \in Q \text{ and } u_1 \in U \quad (7)$$

$$f(u_1, v_1) = \min\{f(u, v) : u \in U \text{ and } v \in Q\} \quad (8)$$

- 2) Move the vertex v_1 from Q into the set U resulting in the new U and Q after the visiting of the vertex v_1 .
- 3) Include the vertex (u_1, v_1) in A , which we record with $\text{tree}[v_1] = u_1$.
- 4) If $Q \neq \emptyset$, go to step II.1. Otherwise, stop.

It is clear that the inductive Procedure 1 terminates after $n - 2$ iterations of the inductive step. After the termination of the procedure, in A are contained $n - 1$ edges, and each vertex of the network N is incident with at least one edge of A .

The following theorem holds.

Theorem 1. *The set A that is constructed by the Procedure 1 is a minimum spanning tree.*

Proof: In the proof we follow [9]. Since the empty set is a subset of every spanning tree, we can assume that in the beginning of Procedure 1 the set A is a subset of some minimum spanning tree T . In the execution of step II.1 of the procedure, we find an edge (u_1, v_1) , such that it satisfies (7) and (8), which means that the edge (u_1, v_1) is a *light edge* that crosses the *cut* (U, Q) . By the definition of the set U , it contains all endpoints of the edges in A , and therefore (u_1, v_1) is not contained in A . Then, from Theorem 23.1 from [9] it follows that there exists a minimum spanning tree T' that contains the set $A \cup \{(u_1, v_1)\}$. It is clear that the set A expanded in this way does not contain cycles because it is a subset of a tree. After the termination of the Procedure 1 the set A contains $n - 1$ edges and therefore it coincides with the minimum spanning tree T' in which it is contained. ■

Besides Theorem 1, each minimum spanning tree can be constructed using an appropriate implementation of the Procedure 1. The following theorem holds.

Theorem 2. *Let T_0 be a minimum spanning tree. Then, the tree T_0 can be constructed using the Procedure 1.*

Proof: Without loss of generality, we assume that the root r is the vertex 1.

I. Base step. We calculate:

$$d = \min\{f(1, v) : v \in G.\text{Adj}[1]\}, \quad (9)$$

$$d_0 = \min\{f(1, v) : (1, v) \in T_0\}. \quad (10)$$

Apparently, $d \leq d_0$. We will prove that $d = d_0$. For the purpose of contradiction, we assume that $d < d_0$.

With v and v_0 we denote two vertices for which the equalities are fulfilled: $f(1, v) = d$, $f(1, v_0) = d_0$, $(1, v_0) \in T_0$ and $v \in N.\text{Adj}[1]$. Then, in the minimum spanning tree T_0 there is a unique path α that has end vertices v_0 and v . Therefore, $\beta = \{(1, v_0)\} \cup \alpha \cup \{(1, v)\}$ is a cycle, and

$T_1 = (T_0 \setminus \{(1, v_0)\}) \cup \{(1, v)\}$ is a spanning tree. It is calculated directly that

$$x(T_0) - x(T_1) = f(1, v_0) - f(1, v) = d_0 - d > 0.$$

Then $x(T_0) > x(T_1)$ which contradicts the condition that T_0 is a minimum spanning tree. Therefore, $d = d_0$. This allows Procedure 1 to transfer the vertex v_0 from Q into U and to define $A = \{(1, v_0)\}$. The cut (U, V) respects the set A that is constructed by the above steps, and A is contained in the minimum spanning tree T_0 .

II. Inductive step. We assume that the set $A \subseteq T_0$ is constructed using Procedure 1 and that the cut (U, Q) is defined. The set U contains the vertices that are endpoints of the edges that belong to A and $Q = V \setminus U$ is the set of vertices that are not yet visited. Besides that, we will assume that $Q \neq \emptyset$, because otherwise the theorem is proved.

We calculate

$$d_1 = \min\{f(u, v) : u \in U \text{ and } v \in Q\}, \quad (11)$$

$$d_0 = \min\{f(u, v) : u \in U, v \in Q \text{ and } (u, v) \in T_0\}. \quad (12)$$

For $i \in \{0, 1\}$ we select vertices u_i and v_i , such that $u_i \in U$, $v_i \in V$, $f(u_i, v_i) = d_i$ and $(u_0, v_0) \in T_0$. Apparently, $d_1 \leq d_0$. We will prove that $d_1 = d_0$. For the purpose of contradiction we assume that $d_1 < d_0$.

Since T_0 is a spanning tree, in T_0 exists a single path γ with end vertices u_1 and v_1 . Then there exists at least one edge $(u, v) \in \gamma$, such that $u \in U$ and $v \in V$. From (12) it follows that $d_0 \leq f(u, v)$. From the assumption $d_1 < d_0$ it follows that:

$$f(u_1, v_1) = d_1 < d_0 \leq f(u, v). \quad (13)$$

It is clear that $\delta = \gamma \cup \{(u_1, v_1)\}$ is a cycle and $T_1 = (T_0 \setminus \{(u, v)\}) \cup \{(u_1, v_1)\}$ is a spanning tree. Therefore, $d_1 = d_0$. This allows Procedure 1 to include into A exactly the edge (u_0, v_0) . Following Procedure 1, we move the vertex v_0 from Q into U . Then it is obvious that the new set A is a subset of T_0 , the set U contains the vertices of the edges that belong to A , and the number of the vertices that are not visited is decremented with one. After $n - 2$ iterations of the inductive step of Procedure 1 we get the set A that coincides with T_0 and the theorem is proved. ■

Theorem 2 allows us to complement Procedure 1 in such a way that, simultaneously with the construction of the minimum spanning tree A , we can also form the list of subproblems with the help of which we can obtain all the remaining minimum spanning trees. We present this extension using Procedure 2. For this purpose, we use the fact that Procedure 1 builds the minimum spanning tree by processing in a standard way the current states of the sets A and Q . This allows us to preserve all possibilities for completing A to a minimum spanning tree that are different from the currently implemented completion. We will store the discovered possibilities in the list S . Such possibilities may appear both in the base step of the procedure and in the inductive step of the procedure.

In Procedure 2 we are not going to use explicitly the set A , but we are going to use its record using the vector $tree$. Besides that, instead of the explicit notation of the sets U and Q , we will store them using the n -components vector $bypassed$.

$$bypassed[v] = \begin{cases} false, & \text{if } v \in Q \\ true, & \text{if } v \in U \end{cases} \quad (14)$$

In this way, we will store each subproblem with the pair $(bypassed, tree)$.

When Procedure 2 starts, the vector $tree$ is the n -component zero vector. Since the vertex r is chosen for the root, all elements of the vector $bypassed$ are $false$, except $bypassed[r]$ which is set to $true$. Thus, the initial problem is stored with the so defined pair of vectors $(bypassed, tree)$, and also the list S is initialized with the empty set.

Procedure 2. *Inductive procedure to construct a minimum spanning tree and all subproblems that define the remaining minimum spanning trees.*

I. Base step.

- 1) For each vertex v_1 , for which

$$f(r, v_1) = \min\{f(r, v) : v \neq r\} \quad (15)$$

define:

$$b_1 = bypassed \text{ and } b_1[v_1] = true, \quad (16)$$

$$t_1 = tree \text{ and } t_1[v_1] = r, \quad (17)$$

$$S = \{(b_1, t_1)\} \cup S. \quad (18)$$

- 2) We choose one element from S and we denote it with $(bypassed, tree)$. This is the subproblem which the procedure will use in its calculations. We remove $(bypassed, tree)$ from S and proceed to the next step.
- 3) If at least one of the components of $bypassed$ is $false$, go to the inductive step. Otherwise, stop.

II. Inductive step.

- 1) For each pair of vertices u_1 and v_1 that satisfy the conditions:

$$bypassed[u_1] = true \text{ and } bypassed[v_1] = false, \quad (19)$$

$$f(u_1, v_1) = \min\{f(u, v) : bypassed[u] = true \text{ and } bypassed[v] = false\}, \quad (20)$$

define

$$b_1 = bypassed \text{ and } b_1[v_1] = true, \quad (21)$$

$$t_1 = tree \text{ and } t_1[v_1] = u_1, \quad (22)$$

$$S = \{(b_1, t_1)\} \cup S. \quad (23)$$

- 2) We denote the latest subproblem (b_1, t_1) that is included in S with $(bypassed, tree)$ and the calculations of the procedure will continue with it. We remove $(bypassed, tree)$ from S and proceed to the next step.

- 3) If at least one of the components of the vector $bypassed$ is $false$, go back for next iteration in step II.1. Otherwise, stop.

We will follow the calculations of Procedure 2.

In step I.1 we discover all vertices v_1 , for which the condition (15) is fulfilled. Each such v_1 satisfies the step I.1 of Procedure 1. With equality (16) we define the vector b_1 that stores the new sets U and Q , if the vertex v_1 is visited, which implements step I.2 of Procedure 1. Also, with equality (17) we record the fact that by traversing the vertex v_1 the tree A being built is augmented with the edge (r, v_1) , which implements step I.3 of Procedure 1. In equality (18), we store in the list S all the subproblems that are obtainable by the base step of Procedure 1.

In step I.2 of Procedure 2 we separate from S the subproblem with which the computations continue.

In step II.1 of Procedure 2 we discover all the pairs of vertices u_1, v_1 that satisfy the conditions of step II.1 of Procedure 1. Also, in equality (21) we store the new sets U and Q , which are defined in step II.2 of Procedure 1. In t_1 , defined by (22), we store the new set A , that would be the result of step II.3 of Procedure 1. We store all subproblems found in this way in the list S with equality (23).

In step II.2 of Procedure 2 we select from S a subproblem that does not violate the logic of Procedure 1.

The inductive step of Procedure 2 is repeated until the main tree $tree$ is constructed. We will emphasize that in such a way the main tree $tree$ is built by implementing Procedure 1. According to Theorem 1, the tree $tree$ is a minimum spanning tree.

The following theorem holds.

Theorem 3. *Let tree and S are constructed by Procedure 2. Then the following two statements hold.*

- 1) The vector $tree$ defines one minimum spanning tree A .
- 2) Each minimum spanning tree T that is different from A is represented by a subproblem stored in S using the inductive step of Procedure 2.

Proof: The proof of the theorem follows from Theorem 1 and Theorem 2. Above we have discussed the second statement of the theorem. We only note that applying the inductive step to a subproblem from S computes the next minimum spanning tree and simultaneously completes the set of subproblems S . ■

Theorem 3 allows us to solve Problem 1 by following Procedure 2. Since Procedure 2 is an extension of the Prim's algorithm [2], as in the case of Dijkstra's algorithm [3], we can adopt the implementation of the abstract data type min-priority queue based on Fibonacci heap data structure [10] to significantly speed up the running time of the algorithm. In the case of the Prim's algorithm implemented with Fibonacci heap, the running time achieved is $O(m + n \lg n)$ [11]. In our case, Fibonacci heap is used to implement the min-priority queue Q that stores the set of vertices that is not yet traversed by the algorithm. The vertices that are stored in Q are keyed

by an attribute d , which we will represent as an n -component vector. In the initialization of the algorithm, the components of the vector d are set to ∞ . At that stage, the set U of the traversed vertices is the empty set. Each change of the set U leads to a change in the min-priority queue Q together with the components of the vector d , for which it holds:

$$d[v] = \begin{cases} \min\{f(u, v) : u \in U\}, & \text{if } v \in Q \\ \infty, & \text{otherwise} \end{cases} \quad (24)$$

In this way, for a vertex v that is not yet traversed, $d[v]$ is the distance of v to the set of traversed vertices.

The implementation of Procedure 2 manages also the attribute vector p with n components, which stores for each vertex $v \in Q$ all the vertices that are in U for which $f(u, v)$ equals the estimated distance of v to the set U .

$$p[v] = \begin{cases} \{u : u \in U, f(u, v) = d[v]\}, & \text{if } v \in Q, d[v] < \infty \\ \{\emptyset\}, & \text{otherwise} \end{cases}$$

We will store each subproblem with the ordered six-tuple:

$$(Q, tree, p, d, bypassed, v), \quad (25)$$

where v is the vertex that was last visited.

In this paper we assume that the network N is given with adjacency lists $N.Adj$, and one of the vertices is selected as the root of the resulting trees. We define a helper function $START(N.Adj, r)$ that takes as input the adjacency lists $N.Adj$ and the selected vertex r for the root of the minimum spanning trees. The function returns the record of the initial problem and the selected root r with the ordered six-tuple of the type (25), where $Q = V \setminus \{r\}$ is min-priority queue keyed by the d attribute vector of the vertices that has components

$$d[w] = \begin{cases} f(r, w), & \text{if } w \in N.Adj[r] \\ \infty, & \text{otherwise} \end{cases} \quad (26)$$

Also, $tree$ is the zero vector, $v = r$, and

$$p[w] = \begin{cases} \{r\}, & \text{if } w \in N.Adj[r] \\ \{\emptyset\}, & \text{otherwise} \end{cases} \quad (27)$$

$$bypassed[w] = \begin{cases} true, & \text{if } w = r \\ false, & \text{if } w \neq r \end{cases} \quad (28)$$

Function $START(N.Adj, r)$ can be implemented, following Algorithm 1 that correctly stores the initial problem with running time complexity that cannot exceed $O(n + m)$.

Example 1. We denote with N_1 the network given on Fig. 1. It is composed of 5 vertices and 10 edges. On the figure, the weights of each edge are given as ordered pair next to it, where the first element is the length and the second is the risk. The network is defined with the following adjacency lists:

$$\begin{aligned} N_1.Adj = [& ((2, 7, 10), (3, 7, 6), (4, 9, 4), (5, 15, 8)), \\ & ((1, 7, 10), (3, 15, 4), (4, 7, 8), (5, 9, 6)), \\ & ((1, 7, 6), (2, 15, 4), (4, 7, 8), (5, 9, 4)), \\ & ((1, 9, 4), (2, 7, 8), (3, 7, 8), (5, 9, 6)), \\ & ((1, 15, 8), (2, 9, 6), (3, 9, 4), (4, 9, 6))]. \end{aligned} \quad (29)$$

Algorithm 1 Function $START(N.Adj, r)$

Input: adjacency lists $N.Adj$ and root vertex r

Output: the record of the initial problem

```

1:  $Q \leftarrow V \setminus \{r\}$ 
2: for  $i \leftarrow 1$  to  $|N.V|$  do
3:    $d[i] \leftarrow \infty$ 
4:    $p[i] \leftarrow \{\emptyset\}$ 
5:    $tree[i] \leftarrow 0$ 
6:   if  $i \neq r$  then
7:      $bypassed[i] \leftarrow false$ 
8:   else
9:      $bypassed[i] \leftarrow true$ 
10:  end if
11: end for
12: for  $i \leftarrow 1$  to  $|N.Adj[r]|$  do
13:    $(w, y, z) \leftarrow N.Adj[r][i]$ 
14:    $d[w] \leftarrow y$ 
15:    $p[w] \leftarrow \{r\}$ 
16: end for
17: return  $(Q, tree, p, d, bypassed, r)$ 

```




Fig. 1. Example network N_1 composed of five vertices with length and risk of each edge given next to it

We choose for root the vertex $r = 1$, and find the record of the initial problem.

Solution. Algorithm 1 determines:

$$\begin{aligned} Q &= \{2, 3, 4, 5\}, tree = (0, 0, 0, 0, 0), \\ p &= (\{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{\emptyset\}), \\ d &= (\infty, \infty, \infty, \infty, \infty), \\ bypassed &= (true, false, false, false, false). \end{aligned}$$

The adjacency list of the root vertex is $N.Adj[r] = ((2, 7, 10), (3, 7, 6), (4, 9, 4), (5, 15, 8))$.

The second **for** loop edits the vectors d and p . We get: $d = (\infty, 7, 7, 9, 15)$ and $p = (\{\emptyset\}, \{1\}, \{1\}, \{1\}, \{1\})$. As a result, the function $START(N_1.Adj, r)$ stores the initial problem for the selected root $r = 1$ as the ordered six-tuple:

$$X = (Q, tree, p, d, bypassed, v), \quad (30)$$

Algorithm 2 Function BRANCHING(V_1, S_1)**Input:** element $V_1 \in D$ and queue S_1 **Output:** the modified queue S_1

```

1:  $(v_1, Q_1) \leftarrow V_1$ 
2:  $b_1 \leftarrow \text{bypassed}, b_1[v_1] \leftarrow \text{true}$ 
3:  $d_1 \leftarrow d, d_1[v_1] \leftarrow \infty$ 
4:  $p_1 \leftarrow p, p_1[v_1] \leftarrow \{\emptyset\}$ 
5: for  $i \leftarrow 1$  to  $|N.Adj[v_1]|$  do
6:    $(w, l, \cdot) \leftarrow N.Adj[v_1][i]$ 
7:   if  $b_1[w] = \text{false}$  then
8:     if  $d_1[w] > l$  then
9:        $d_1[w] \leftarrow l, p_1[w] \leftarrow \{v_1\}$ 
10:    else if  $d_1[w] = l$  then
11:       $p_1[w] \leftarrow p_1[w] \cup \{v_1\}$ 
12:    end if
13:  end if
14: end for
15:  $U \leftarrow p[v_1]$   $\{u \in U \text{ exactly when } u \notin Q$ 
    $\text{and } f(u, v_1) = d[v_1]\}$ 
16: for  $i \leftarrow 1$  to  $|U|$  do
17:    $u \leftarrow U[i]$ 
18:    $t_1 \leftarrow \text{tree}, t_1[v_1] \leftarrow u$ 
19:    $S_1 \leftarrow S_1 \cup \{(Q_1, t_1, p_1, d_1, b_1, v_1)\}$ 
20: end for
21: return  $S_1$ 

```

where

$$\begin{aligned}
Q &= \{2, 3, 4, 5\}, \text{tree} = (0, 0, 0, 0, 0), \\
p &= (\{\emptyset\}, \{1\}, \{1\}, \{1\}, \{1\}), d = (\infty, 7, 7, 9, 15), \\
\text{bypassed} &= (\text{true}, \text{false}, \text{false}, \text{false}, \text{false}), v = 1.
\end{aligned}$$

To implement Procedure 2 we will use the following three helper functions: EXTRACT(Q), BRANCHING(V_1, S_1) and OPEN-PRIM(X, S), with which we will analyze the current subproblem $(Q, \text{tree}, p, d, \text{bypassed}, v)$.

The function EXTRACT(Q) discovers all elements v_j in the priority queue Q for which $d[v_j] = \min\{d[v] : v \in Q\}$ and composes a deque (double-ended queue) D with elements (v_j, Q_j) , where $Q_j = Q \setminus \{v_j\}$. In the beginning of D are stored the elements for which $v \in p[v_j]$.

To each element V_1 of the deque D we apply the helper function BRANCHING(V_1, S_1). It adds to a predefined queue S_1 all subproblems that can complete the current tree tree if an element V_1 of D is selected. The function BRANCHING(V_1, S_1) can be implemented following Algorithm 2, in which we suppose that the queue S_1 is defined, and an element $V_1 \in D$ is chosen.

We will analyze the execution of Algorithm 2. In lines 2 and 3 the algorithm marks the vertex v_1 as traversed. Then it calculates the new values of d_1 and p_1 . Initially, it records that v_1 is not an element of Q . Then it changes the values of the estimated distance d and list of parents p when v_1 becomes a traversed vertex. For the resulting d_1 and p_1 the following equalities hold.

Algorithm 3 Function OPEN-PRIM(X, S)**Input:** subproblem X and the stack of subproblems S **Output:** minimum spanning tree tree and modified S

```

1:  $(Q, \text{tree}, p, d, \text{bypassed}, v) \leftarrow X$ 
2: while  $Q \neq \emptyset$  do
3:    $D \leftarrow \text{EXTRACT}(Q)$ 
4:   let  $S_1$  be an empty queue
5:   while  $D \neq \emptyset$  do
6:      $V_1 \leftarrow \text{FRONT}(D)$ 
7:      $\text{POPFRONT}(D)$ 
8:      $S_1 \leftarrow \text{BRANCHING}(V_1, S_1)$ 
9:   end while
10:   $S \leftarrow S_1 \cup S$   $\{\text{preserve the order of elements in } S_1\}$ 
11:   $(Q, \text{tree}, p, d, \text{bypassed}, v) \leftarrow S.\text{top}$ 
12:   $\text{POP}(S)$ 
13: end while
14: return  $\text{tree}, S$ 

```

1) If w belongs to the set $Q_1 \cap N.Adj(v_1)$, then

$$d_1[w] = \begin{cases} d[w], & \text{if } f(v_1, w) = d[w] \\ f(v, w), & \text{if } f(v_1, w) < d[w] \end{cases} \quad (31)$$

$$p_1[w] = \begin{cases} p[w] \cup \{v_1\}, & \text{if } f(v_1, w) = d[w] \\ \{v_1\}, & \text{if } f(v_1, w) < d[w] \\ p[w], & \text{if } f(v_1, w) > d[w] \end{cases} \quad (32)$$

2) If w belongs to the set $(Q_1 \setminus N.Adj[v_1]) \cup U$, then $d_1[w] = d[w]$ and $p_1[w] = p[w]$.3) $d_1[v_1] = \infty$ and $p_1[v_1] = \{\emptyset\}$.

Therefore, Algorithm 2 correctly separates all resulting subproblems, if the element V_1 is selected from the deque D .

Using functions EXTRACT(Q) and BRANCHING(V_1, S_1) we define the function OPEN-PRIM(X, S) in Algorithm 3 that implements the inductive step of Procedure 2. We assume that are defined the subproblem $X = (Q, \text{tree}, p, d, \text{bypassed}, v)$ with $Q \neq \emptyset$ and the stack S of subproblems of the initial problem. The input of the function OPEN-PRIM(X, S) is the subproblem X and the stack S , and its output is a minimum spanning tree stored in the vector tree and the modified stack S . During its execution the function can push new subproblems into S .

Algorithm 3 correctly implements the inductive step of Procedure 2, more precisely, lines from 3 to 10 implement step II.1, and lines 11, 12 implement step II.2.

Indeed, the function EXTRACT(Q) in line 3 of the algorithm finds all vertices v_1 for which there exists such vertex u_1 for which Eq. (19) and (20) hold. In the inner **while** loop the algorithm determines the queue S_1 of all subproblems that can complete the tree tree if v_1 is chosen to be traversed. In this case, we use the correctness of the function BRANCHING(V_1, S_1). This, in particular, guarantees the fulfillment of Eq. (19), (20), (21) and (22). With line 10, the algorithm stores the elements of S_1 in the top of the stack S , preserving their

order, and thus the implementation of step II.2 of Procedure 2 is completed.

The running time of Algorithm 3 depends on the efficiency of the functions of the min-priority queue Q abstract data type. In the case in which Fibonacci heap data structure is used for the implementation of Q , the running time complexity of Algorithm 3 is $O(m + n \lg n)$ (see [11] and also [9]). This is true, because the proposed algorithm is an extension of the Prim's algorithm [2] and differs from it in a constant factor.

We note that the computations in step I.1 of Procedure 2 are a special case of the computations in step II.1. This enables the computations in step 1 to be implemented with the OPEN-PRIM(X, S) function as well.

We will illustrate the above with the following Example 2. In Example 1 we already proved that Problem 1 for the network N_1 when $r = 1$ is written with the ordered six-tuple X defined with Eq. (30). We denote by S an empty stack.

Example 2. For the subproblem X and stack S defined above, we will prove that the first iteration of the outer while loop of Algorithm 3 implements the base step of Procedure 2.

Solution. From the definition of X given in Eq. (30) we know that $Q = [2, 3, 4, 5]$ and $d = (\infty, 7, 7, 9, 15)$. Then, the function EXTRACT(Q) defines the deque $D = [\langle 3, (2, 4, 5) \rangle, \langle 2, (3, 4, 5) \rangle]$.

In this case, the inner **while** loop of Algorithm 3 executes two iterations. Each iteration, using the BRANCHING(V_1, S_1) function, defines a subproblem with which the current tree stored in the vector $tree$ can be completed. The first iteration defines the subproblem

$$X_1 = (Q_1, t_1, p_1, d_1, b_1, 3), \quad (33)$$

where

$$Q_1 = [2, 4, 5], t_1 = (0, 0, 1, 0, 0),$$

$$p_1 = (\{\emptyset\}, \{1\}, \{\emptyset\}, \{3\}, \{3\}),$$

$$d_1 = (\infty, 7, \infty, 7, 9), b_1 = (true, false, true, false, false).$$

The second iteration defines the subproblem

$$X_2 = (Q_2, t_2, p_2, d_2, b_2, 2), \quad (34)$$

where

$$Q_2 = [3, 4, 5], t_2 = (0, 1, 0, 0, 0),$$

$$p_2 = (\{\emptyset\}, \{\emptyset\}, \{1\}, \{2\}, \{2\}),$$

$$d_2 = (\infty, \infty, 7, 7, 9), b_2 = (true, true, false, false, false).$$

In this way, step I.1 of Procedure 2 is implemented, which is verified directly. Lines 10, 11 and 12 of Algorithm 3 implement step I.2 of Procedure 2. Thus, the first iteration of the outer **while** loop of the algorithm implements the base step of the procedure. Moreover, following the depth-first search principle, subproblem X_1 remains as the current problem, and only subproblem X_2 remains in the stack S . ■

Example 2 clearly demonstrates that Algorithm 3 implements Procedure 2. In this case, the outer **while** loop will perform $n - 1$ iterations. The first iteration will implement the

Algorithm 4 Function SUBPROBLEMS(N)

Input: the network N

Output: minimum spanning tree $tree$ and subproblems S

- 1: let S be an empty stack
 - 2: $X \leftarrow \text{START}(N.Adj, r)$
 - 3: $(tree, S) \leftarrow \text{OPEN-PRIM}(X, S)$
 - 4: **return** $tree, S$
-

initial step of the procedure, and the remaining $n - 2$ iterations will implement the inductive steps. This allows the solution of Problem 1 to be obtained with Algorithm 4.

Example 3. We will find a single minimum spanning tree and compile a list S of the subproblems that can be used to find the remaining minimum spanning trees of the network N_1 from Example 1 using Algorithm 4.

Solution. In lines 1 and 2 the algorithm defines an empty stack S and stores the initial problem in the ordered six-tuple $X = (Q, tree, p, d, bypassed, 1)$. In Example 1 we have proved that X is defined with equality (30).

The computations in line 3 are implemented using Algorithm 3. We will follow these calculations for the given case.

In Example 2, we found that the first iteration of the outer **while** loop implements the base step of Procedure 2 and defines as the current subproblem X_1 , which is given by equality (33). Additionally, the S stack is also defined, and it contains only the subproblem X_2 which is given by (34).

Since the current $Q = [2, 4, 5]$, the outer **while** loop executes its second iteration, and thus we execute the first iteration of the inductive step of Procedure 2. At that stage, the attribute vector d , that is used as min-priority queue keys, stores $d = (\infty, 7, \infty, 7, 9)$. The function EXTRACT(Q) determines that $D = [\langle 4, (2, 5) \rangle, \langle 2, (4, 5) \rangle]$.

The inner **while** loop of Algorithm 3 executes two iterations. On each iteration, using the function BRANCHING(V_1, S_1), it defines a subproblem with which the tree stored in $tree$ can be completed. In line 10 these two subproblems are stored on the top of the stack S . We get $S = [X_{11}, X_{12}, X_2]$, where the subproblem X_2 is defined by (34), and

$$X_{11} = ((2, 5), (0, 0, 1, 3, 0), (\{\emptyset\}, \{1, 4\}, \{\emptyset\}, \{\emptyset\}, \{3, 4\}), (\infty, 7, \infty, \infty, 9), (true, false, true, true, false), 4), \quad (35)$$

$$X_{12} = ((4, 5), (0, 1, 1, 0, 0), (\{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{3, 2\}, \{3, 2\}), (\infty, \infty, \infty, 7, 9), (true, true, true, false, false), 2). \quad (36)$$

In line 11, we set the new current subproblem to be X_{11} , and then in line 12 we pop from S the top element. Now, the new queue Q has two elements and $Q = [2, 5]$. Therefore, the outer while loop implements a third iteration and simultaneously starts the execution of the next iteration of the inductive step of Procedure 2. After the execution of the third iteration, line 11 defines the new

current subproblem with $Q = [5]$, $tree = (0, 1, 1, 3, 0)$, $p = (\{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{3, 4, 2\})$, $d = (\infty, \infty, \infty, \infty, 9)$, $bypassed = (true, true, true, true, false)$ and $v = 2$.

Besides that, in line 12, the new stack S is defined, which differs from the stack S defined in the previous iteration in that the subproblem is pushed:

$$\begin{aligned} X_{112} = & ((5), (0, 4, 1, 3, 0), (\{\emptyset\}, \{\emptyset\}, \\ & \{\emptyset\}, \{\emptyset\}, \{3, 4, 2\})(\infty, \infty, \infty, \infty, 9), \\ & (true, true, true, true, false), 2). \end{aligned} \quad (37)$$

In this way, $S = [X_{112}, X_{12}, X_2]$.

Since the current $Q \neq \emptyset$, the outer loop of the algorithm executes its fourth iteration and thus the next iteration of the inductive step of Procedure 2. This time, in line 11 the current $Q = \{\emptyset\}$, $tree = (0, 1, 1, 3, 3)$, every element of p is the empty set, every element of d is infinity, every element of $bypassed$ is $true$ and $v = 1$.

Besides that, in the stack S of the previous iteration are pushed two new subproblems:

$$\begin{aligned} X_{1112} = & ((\emptyset), (0, 1, 1, 3, 4), (\{\emptyset\}, \{\emptyset\}, \\ & \{\emptyset\}, \{\emptyset\}, \{\emptyset\})(\infty, \infty, \infty, \infty, \infty), \\ & (true, true, true, true, true), 5), \end{aligned} \quad (38)$$

$$\begin{aligned} X_{1113} = & ((\emptyset), (0, 1, 1, 3, 2), (\{\emptyset\}, \{\emptyset\}, \\ & \{\emptyset\}, \{\emptyset\}, \{\emptyset\})(\infty, \infty, \infty, \infty, \infty), \\ & (true, true, true, true, true), 5). \end{aligned} \quad (39)$$

At the end of the fourth iteration, the queue Q is the empty set. Therefore, no further iterations of the outer **while** loop of the algorithm are executed and, moreover, the iterations of the inductive step of Procedure 2 also stop. In this way, the calculations in line 3 of Algorithm 4 stop and it returns:

$$\begin{aligned} tree = & (0, 1, 1, 3, 3) \text{ and} \\ S = & [X_{1112}, X_{1113}, X_{112}, X_{12}, X_2], \end{aligned} \quad (40)$$

where the elements of the stack S are defined by Eq. (38), (39), (37), (36) and (34).

From Theorem 3 it follows that the tree $tree = (0, 1, 1, 3, 3)$ is a minimum spanning tree. It can be directly computed that the weight according to the length objective function of that tree is 30, and it can be stored as the set of edges $A = \{(1, 2), (1, 3), (3, 4), (3, 5)\}$.

Also, from Theorem 3 follows that that any other minimum spanning tree can be obtained from the stack of subproblems S . In the particular case, directly from Eq. (38) and (39) we establish that the first subproblem defines a tree $(0, 1, 1, 3, 4)$ and the second subproblem defines a tree $(0, 1, 1, 3, 2)$. We can directly verify that the two new spanning trees have length 30 again. ■

We will note that the correctness of Algorithm 4 follows from the correctness of Algorithm 3 and the fact that the first iteration of the outer **while** loop of Algorithm 3 implements the initial step of Procedure 2. Moreover, the running time complexity of Algorithm 4 coincides with the complexity of Algorithm 3.

Corollary 1. *Algorithm 4 can be edited in such a way that it returns a list of all minimum spanning trees.*

For example, using the algorithm from Corollary 1 we get that the network N_1 from Example 1 has exactly 12 minimum spanning trees. These are the trees given in the list S :

$$\begin{aligned} S = & \{(0, 1, 1, 3, 3), (0, 1, 1, 3, 4), (0, 1, 1, 3, 2), \\ & (0, 4, 1, 3, 3), (0, 4, 1, 3, 4), (0, 4, 1, 3, 2), \\ & (0, 1, 1, 2, 3), (0, 1, 1, 2, 2), (0, 1, 1, 2, 4), \\ & (0, 1, 4, 2, 2), (0, 1, 4, 2, 4), (0, 1, 4, 2, 3)\}. \end{aligned} \quad (41)$$

IV. COMPLETE PARETO FRONT ALGORITHM

In this section we will give the solution of the main problem considered in this paper.

Problem 2 (Main problem). *Let N be a connected network given with its adjacency lists $N.Adj$. Compose a list P of all classes of equivalent, Pareto optimal trees.*

We will solve Problem 2 using the following Procedure 3.

Procedure 3. *Compose a list of all classes of equivalent, Pareto optimal trees.*

- 1) Calculate the minimum length l of a spanning tree, $l = \min\{x(t) : t \in W\}$.
- 2) Calculate the minimum risk r for the minimum spanning trees, $r = \min\{y(t) : t \in W \text{ and } x(t) = l\}$.
- 3) Define the set T of all spanning trees t for which $x(t) = l$ and $y(t) = r$.
- 4) Define the sets $P = P \cup T$ and $W_1 = \{t : t \in W \text{ and } y(t) < r\}$.
- 5) If $W_1 \neq \emptyset$, set $W = W_1$ and go to step 1. Otherwise, stop.

After the termination of Procedure 3 in P are stored all equivalent, Pareto optimal trees.

The following lemma holds.

Lemma 1. *Procedure 3 correctly computes the list P of the classes of equivalent Pareto optimal trees of the network N .*

Proof: It is sufficient to prove that the invariant given in Procedure 3 correctly separates the consecutive class of equivalent Pareto optimal trees.

With steps 1, 2 and 3 we define:

$$\begin{aligned} l_1 = & \max\{x(t) : t \in W\}, \\ X_1 = & \{t : t \in W \text{ and } x(t) = l_1\}, \\ r_1 = & \max\{y(t) : t \in X_1\}, \\ T_1 = & \{t : t \in X_1 \text{ and } y(t) = r_1\}. \end{aligned} \quad (42)$$

We denote $W_1 = \{t : t \in W \text{ and } y(t) < r\}$ and $Z_1 = W \setminus (W_1 \cup T_1)$.

Let t_1 be an arbitrary tree from the set T_1 . If $Z_1 \neq \emptyset$, then for each $t \in Z_1$ one of the cases (43) or (44) holds.

$$x(t) \geq x(t_1) \text{ and } y(t) > y(t_1) \quad (43)$$

$$x(t) > x(t_1) \text{ and } y(t) = y(t_1) \quad (44)$$

Therefore, the tree t is dominated by the tree t_1 , written $t \prec t_1$.

If $W_1 \neq \emptyset$, then for each $t \in W_1$ the inequalities (45) hold.

$$x(t) > x(t_1) \text{ and } y(t) < y(t_1) \quad (45)$$

Therefore, in this case, t and t_1 are not comparable.

As a result, the first iteration of Procedure 3 partitions the set of all minimum spanning trees W into three disjoint subsets T_1 , Z_1 , and W_1 .

If W_1 is the empty set, then Procedure 3 stops and the problem has single class of equivalent, Pareto optimal trees: the class T_1 .

If W_1 is not the empty set, then again T_1 is a class of equivalent, Pareto optimal trees, because the elements of W_1 and T_1 are not comparable. Then Procedure 3 returns to step 1 for the next iteration. The next iteration partitions the set W_1 into three disjoint subsets T_2 , Z_2 , and W_2 . In this case:

- T_2 is a class of equivalent, Pareto optimal trees;
- Z_2 contains the elements of W_1 that are dominated by the elements of T_2 ;
- W_2 contains the elements of W_1 that are not comparable to the elements of T_2 .

Since W is a finite set, the procedure stops after finite number of iterations. At the end of the proof, we will emphasize that the procedure makes exactly as many iterations as there are different classes of equivalent, Pareto optimal trees. ■

In the implementation of Procedure 3, besides helper functions $\text{START}(N.Adj, r)$ and $\text{OPEN-PRIM}(X, S)$ given in Sec. III, we will use the helper functions $\text{CONNECTED}(N.Adj)$, $\text{RESTRICT}(N.Adj, c)$, $\text{NEW}(t, T, c)$ and $\text{MPOT}(N.Adj, r)$.

The function $\text{CONNECTED}(N.Adj)$ is a predicate function that performs depth-first search in the network, and returns *true* if N is connected, and *false* if N is not connected.

The function $\text{RESTRICT}(N.Adj, c)$ separates a subnetwork from the network N . The subnetwork contains only those edges of N that have risk strictly less than c . The function returns the adjacency lists of the separated subnetwork.

The function $\text{NEW}(t, T, c)$ separates the trees with minimum risk. We assume that T is a list of trees that have risk c . We further assume that c is the minimum currently detected risk. The function $\text{NEW}(t, T, c)$ checks whether t can improve the currently minimum c . If this is the case, then we define $T = \{t\}$ and $c = \text{risk}(t)$. Also, if $\text{risk}(t) = c$ and $t \notin T$, then the function $\text{NEW}(t, T, c)$ also adds to the list T the tree t . The function $\text{NEW}(t, T, c)$ can be implemented based on the Algorithm 5.

Problem 3. Let N be a connected network represented by its adjacency lists $N.Adj$. Compose a list T of all Pareto optimal trees that have minimum length.

The function $\text{MPOT}(N.Adj, r)$ given in Algorithm 6 solves Problem 3.

The correctness of the Algorithm 6 follows directly from Lemma 1 and the correctness of Algorithm 3 and Algorithm 5. It is easy to verify that the running time complexity of Algorithm 6 is evaluated to $O(L(m+n \lg n))$, where L is the

Algorithm 5 Function $\text{NEW}(t, T, c)$

Input: a tree t , list of trees T with risk c
Output: updated list of trees T with risk c

- 1: **if** $\text{risk}(t) < c$ **then**
- 2: $T \leftarrow \{t\}$
- 3: $c \leftarrow \text{risk}(t)$
- 4: **else if** $\text{risk}(t) = c$ and $t \notin T$ **then**
- 5: $T \leftarrow T \cup \{t\}$
- 6: **end if**
- 7: **return** T, c

Algorithm 6 Function $\text{MPOT}(N.Adj, r)$

Input: adjacency lists $N.Adj$ and root vertex r
Output: all minimum length Pareto optimal trees T , risk c

- 1: let T be an empty list, S be an empty stack
- 2: $c \leftarrow \infty$
- 3: $\text{PUSH}(S, \text{START}(N.Adj, r))$
- 4: **while** $S \neq \emptyset$ **do**
- 5: $X \leftarrow S.top$
- 6: $\text{POP}(S)$
- 7: $(t, S_1) \leftarrow \text{OPEN-PRIM}(X, S)$
- 8: $(T, c) \leftarrow \text{NEW}(t, T, c)$
- 9: **while** $S_1 \neq \emptyset$ and $S_1[1, 1] = \emptyset$ **do**
- 10: $(Q, t, p, d, b, v) \leftarrow S_1.top$
- 11: $\text{POP}(S_1)$
- 12: $(T, c) \leftarrow \text{NEW}(t, T, c)$
- 13: **end while**
- 14: $S \leftarrow S_1 \cup S$ {preserve the order of elements in S_1 }
- 15: **end while**
- 16: **return** T, c

number of minimum spanning trees and $O(m+n \lg n)$ is the complexity of the function $\text{OPEN-PRIM}(X, S)$.

Example 4. Compose a list of all minimum spanning trees that are Pareto optimal for the network N_1 in Example 1.

Solution. We will trace the calculations of the outer **while** loop of Algorithm 6.

First iteration of the loop, in lines 5 and 6 stores the initial problem with the six-tuple X defined with Eq. (30) and empties the stack S . In line 7 the call to the function $\text{OPEN-PRIM}(X, S)$ calculates (t, S_1) . In Example 3 we provided a detailed proof that $t = \text{tree}$ and $S_1 = S$, where tree and S are defined by Eq. (40).

After that, in line 8 the function $\text{NEW}(t, T, c)$ edits the current record and we get $c = 10$ and $T = \{(0, 1, 1, 3, 3)\}$.

In lines from 9 to 13, the inner **while** loop performs two iterations and completes the set T . The current record is given in Eq. (46).

$$c = 10, T = \{(0, 1, 1, 3, 3), (0, 1, 1, 3, 4), (0, 1, 1, 3, 2)\} \quad (46)$$

The elements of the stack S_1 are stored on the top of the stack S preserving the order of the elements in S_1 and $S = [X_{112}, X_{12}, X_2]$, where the elements of S are defined

respectively by Eq. (37), (36), and (34). Since $S \neq \emptyset$, the loop proceeds to its second iteration.

Second iteration of the loop applies the function OPEN-PRIM(X, S) to the subproblem X_{112} and the stack $S = [X_{12}, X_2]$. It calculates the minimum spanning tree $t = (0, 4, 1, 3, 3)$ and the stack

$$S_1 = [\langle Q_1, (0, 4, 1, 3, 4), p_1, d_1, b_1, 5 \rangle, \langle Q_2, (0, 4, 1, 3, 2), p_2, d_2, b_2, 5 \rangle],$$

where

$$\begin{aligned} Q_1 = Q_2 = [\emptyset], p_1 = p_2 = (\{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{\emptyset\}, \{\emptyset\}), \\ d_1 = d_2 = (\infty, \infty, \infty, \infty, \infty), \\ b_1 = b_2 = (true, true, true, true, true). \end{aligned}$$

In line 8 the function NEW(t, T, c) improves the current record and we get $c = 8$ and $T = \{(0, 4, 1, 3, 3)\}$.

In lines from 9 to 13, the inner **while** loop again performs two iterations and completes the set T . The resulting current record is:

$$c = 8, T = \{(0, 1, 1, 3, 3), (0, 4, 1, 3, 4), (0, 4, 1, 3, 2)\}. \quad (47)$$

Since the stack $S = [X_{12}, X_2]$ is not empty, the loop proceeds to its next iteration.

The outer **while** loop executes six more iterations. The minimum spanning trees that are discovered by these iterations have a risk greater than $c = 8$. Therefore, the current record does not change. The function MPOT($N.Adj, r$) returns the set T and the risk c which are defined by Eq. (47).

From the correctness of the function OPEN-PRIM(X, S) it follows that the function MPOT($N.Adj, r$) has traversed all minimum spanning trees. From the correctness of the function NEW(t, T, c) it follows that in T are separated the minimum spanning trees that have minimum risk.

It is easy to observe that in such a way the function MPOT($N.Adj, r$) implements the first iteration of Procedure 3. Then, from Lemma 1, in particular, it follows that T is a class of equivalent, Pareto optimal spanning trees. ■

Corollary 2. *Let the network N_2 be obtained from the network N with the risk of each edge changed to 1. Then the function call MPOT($N_2.Adj, r$) composes the list T_1 of all minimum spanning trees of the network N .*

For example, for the network N_1 of Example 1, we get that $T_1 = S$, where S is given by Eq. (41).

Corollary 3. *Let the network N_3 be obtained from the network N with both risk and length of each edge changed to 1. Then the function call MPOT($N_3.Adj, r$) composes the list T_2 of all spanning trees of the network N .*

Using the helper functions defined above, we will solve the main Problem 2. The proposed solution is Algorithm 7.

The list P that results from Algorithm 7 is a solution of the main Problem 2. This follows directly from Lemma 1 and the fact that Algorithm 7 implements Procedure 3.

Indeed, let us denote with W the set of all spanning trees of the network N . The first iteration of the algorithm separates

Algorithm 7 Function CPOT($N.Adj, r$)

Input: adjacency lists $N.Adj$ and root vertex r

Output: all classes of equivalent, Pareto optimal trees P

```

1: let  $P$  be an empty list
2:  $ind \leftarrow true$ 
3: while  $ind = true$  do
4:    $(T, c) \leftarrow MPOT(N.Adj, r)$ 
5:    $P \leftarrow P \cup \{T\}$ 
6:    $N.Adj \leftarrow RESTRICT(N.Adj, c)$ 
7:    $ind \leftarrow CONNECTED(N.Adj)$ 
8: end while
9: return  $P$ 

```

the set T from those minimum spanning trees that have minimum risk. We denote with c the risk and with l the length of an arbitrary tree of T .

In line 5 of Algorithm 7 the set T is included into the list P . Then we denote $W_1 = \{t : t \in W \text{ and } y(t) < c\}$. This implements steps from 1 to 4 of Procedure 3.

In line 6 of Algorithm 7 the function RESTRICT($N.Adj, c$) separates the subnetwork N' that contains only those edges of the network N that have a risk strictly less than c . We note that a tree t belongs to W_1 if and only if it is a spanning tree of the subnetwork N' . Therefore, $W_1 \neq \emptyset$ if and only if the subnetwork N' is connected. This proves that Algorithm 7 will execute next iteration exactly when Procedure 3 executes its next iteration.

The **while** loop terminates when the subnetwork N' is not connected and all K number of classes of Pareto equivalent trees are discovered. Therefore, from the computational complexity of the function MPOT($N.Adj, r$) it follows that Algorithm 7 has running time $O(KL(m + n \lg n))$.

The following Example 5 clarifies the proof of the correctness of Algorithm 7.

Example 5. *We will examine the network N_5 that is composed by 9 vertices and 14 edges, and is defined by the adjacency lists given in (48).*

$$\begin{aligned} N_5.Adj = [\langle (2, 4, 2), (3, 8, 2) \rangle, \\ \langle (1, 4, 2), (3, 11, 6), (4, 8, 2) \rangle, \\ \langle (1, 8, 2), (2, 11, 6), (5, 7, 6), (6, 1, 2) \rangle, \\ \langle (2, 8, 2), (5, 2, 6), (7, 4, 4), (8, 7, 8) \rangle, \\ \langle (3, 7, 6), (4, 2, 6), (6, 6, 4) \rangle, \\ \langle (3, 1, 2), (5, 6, 4), (7, 2, 4) \rangle, \\ \langle (4, 4, 4), (6, 2, 4), (8, 14, 2), (9, 10, 4) \rangle, \\ \langle (4, 7, 8), (7, 14, 2), (9, 9, 8) \rangle, \\ \langle (7, 10, 4), (8, 9, 8) \rangle] \end{aligned} \quad (48)$$

Using Algorithm 7 we will compose a list of all classes of equivalent Pareto optimal trees.

Solution. We set $r = 1$ and denote with W the set of all spanning trees of the network N_5 .

First iteration of the **while** loop of the algorithm, using the function $\text{MPOT}(N.\text{Adj}, r)$ in line 4 calculates

$$T_1 = \{\{(0, 1, 6, 2, 4, 7, 4, 4, 8), (0, 1, 1, 7, 4, 3, 6, 4, 8)\}\} \text{ and } c_1 = 8. \quad (49)$$

We define the sets

$$W_1 = \{t : t \in W \text{ and } y(t) < c_1\}, \\ Z_1 = \{t : t \in W, y(t) \geq c_1 \text{ and } t \notin T_1\}.$$

Then, obviously, the three sets T_1 , Z_1 , and W_1 have no common elements and

$$W = T_1 \cup Z_1 \cup W_1. \quad (50)$$

Let $t_0 \in T_1$ and $t \in Z_1 \cup W_1$. From the correctness of the function $\text{MPOT}(N.\text{Adj}, r)$ it follows that $x(t_0) \leq x(t)$. Moreover, if $t \in Z_1$, then the following two cases are possible:

- $x(t_0) < x(t)$ and $y(t_0) = c_1 = y(t)$;
- $x(t_0) \leq x(t)$ and $y(t_0) = c_1 < y(t)$.

Therefore, $t \prec t_0$.

If $t \in W_1$, then $y(t_0) = c_1 > y(t)$. From the correctness of the function $\text{MPOT}(N.\text{Adj}, r)$ it follows that $x(t_0) < x(t)$. Then, t_0 and t are not comparable. Therefore, every tree $t_0 \in T_1$ is Pareto optimal, and in line 5 of the algorithm T_1 is correctly included in the list P .

We denote with l_1 the length of any tree $t_0 \in T_1$. In the considered example $l_1 = x(t_0) = 37$.

In line 6 the subnetwork N' is defined that contains only those edges of N that have risk strictly less than $c_1 = 8$. We get

$$N'.\text{Adj} = [\langle(2, 4, 2), (3, 8, 2)\rangle, \langle(1, 4, 2), (3, 11, 6), (4, 8, 2)\rangle, \langle(1, 8, 2), (2, 11, 6), (5, 7, 6), (6, 1, 2)\rangle, \langle(2, 8, 2), (5, 2, 6), (7, 4, 4)\rangle, \langle(3, 7, 6), (4, 2, 6), (6, 6, 4)\rangle, \langle(3, 1, 2), (5, 6, 4), (7, 2, 4)\rangle, \langle(4, 4, 4), (6, 2, 4), (8, 14, 2), (9, 10, 4)\rangle, \langle(7, 14, 2)\rangle, \langle(7, 10, 4)\rangle]. \quad (51)$$

We note that W_1 is the set of all spanning trees of the subnetwork N' . This fact can be verified directly by proving that a tree $t \in W_1$ if and only if t is a spanning tree of N' .

Let $t \in W_1$. Then t is a spanning tree of N and, in particular, every vertex of N' is incident to an edge of t . Furthermore, every edge of t , by the definition of W_1 , has a risk strictly less than c_1 . Therefore, t is a spanning tree of N' . Analogously, it is verified that if t is a spanning tree of N' , then $t \in W_1$.

In line 7 it is verified that the network N' is connected, and the loop proceeds to its second iteration. In particular, this means that $W_1 \neq \emptyset$.

Second iteration of the **while** loop, using function call $\text{MPOT}(N'.\text{Adj}, r)$ calculates

$$T_2 = \{\{(0, 1, 6, 2, 4, 7, 4, 7, 7), (0, 1, 1, 7, 4, 3, 6, 7, 7)\}\} \text{ and } c_2 = 6. \quad (52)$$

We denote with l_2 the length of any tree $t_0 \in T_2$. In the considered example, $l_2 = x(t_0) = 45$. Then we define the sets

$$W_2 = \{t : t \in W_1 \text{ and } y(t) < c_2\}, \\ Z_2 = \{t : t \in W_1, y(t) \geq c_2 \text{ and } t \notin T_2\}.$$

Then apparently the three sets T_2 , Z_2 and W_2 have no common elements and

$$W = T_2 \cup Z_2 \cup W_2. \quad (53)$$

Then the equality (50) is written in the form

$$W = T_1 \cup Z_1 \cup T_2 \cup Z_2 \cup W_2. \quad (54)$$

Let $t_0 \in T_2$. We will prove that t_0 is a Pareto optimal tree.

Let t be an arbitrary tree from the set $W \setminus T_2$. Then obviously the following four cases are possible: 1) $t \in W_2$; 2) $t \in Z_2$; 3) $t \in Z_1$ and 4) $t \in T_1$. Repeating the reasoning from the first iteration of the loop, we immediately find the following.

- 1) If $t \in W_2$, then t and t_0 cannot be compared.
- 2) If $t \in Z_2$, then $t \prec t_0$.
- 3) If $t \in Z_1$ two subcases are possible:
 - $x(t) < l_2$ and then t and t_0 cannot be compared;
 - $x(t) \geq l_2$ and then $t \prec t_0$.
- 4) If $t \in T_1$, then t and t_0 cannot be compared because $t_0 \in W_1$.

Therefore, every tree $t_0 \in T_2$ is Pareto optimal, and in line 5 of the algorithm T_2 is correctly included in the list P .

In line 6 the subnetwork N'' is defined that contains only those edges of N that have risk strictly less than $c_2 = 6$. The resulting adjacency lists of the network N'' is:

$$N''.\text{Adj} = [\langle(2, 4, 2), (3, 8, 2)\rangle, \langle(1, 4, 2), (4, 8, 2)\rangle, \langle(1, 8, 2), (6, 1, 2)\rangle, \langle(2, 8, 2), (7, 4, 4)\rangle, \langle(6, 6, 4)\rangle, \langle(3, 1, 2), (5, 6, 4), (7, 2, 4)\rangle, \langle(4, 4, 4), (6, 2, 4), (8, 14, 2), (9, 10, 4)\rangle, \langle(7, 14, 2)\rangle, \langle(7, 10, 4)\rangle]. \quad (55)$$

As above, we find that W_2 is the set of all spanning trees of the subnetwork N'' . In line 7 of the algorithm we find that N'' is connected and the loop proceeds to its third iteration.

Third iteration of the **while** loop calls $\text{MPOT}(N''.\text{Adj}, r)$ and calculates

$$T_3 = \{\{(0, 1, 6, 2, 6, 7, 4, 7, 7), (0, 1, 1, 7, 6, 3, 6, 7, 7)\}\} \text{ and } c_3 = 4. \quad (56)$$

We denote with l_3 the length of any tree $t_0 \in T_3$. In the considered example $l_3 = x(t_0) = 49$. Then we define the sets

$$W_3 = \{t : t \in W_2 \text{ and } y(t) < c_3\}, \\ Z_3 = \{t : t \in W_1, y(t) \geq c_3 \text{ and } t \notin T_3\}.$$

Then, apparently, the three sets T_3 , Z_3 and W_3 have no common elements and

$$W = T_1 \cup Z_1 \cup T_2 \cup Z_2 \cup T_3 \cup Z_3 \cup W_3. \quad (57)$$

Every tree t_0 of T_3 is Pareto optimal. The proof is completely analogous to the proof that T_2 contains only Pareto optimal trees. Therefore, T_3 is correctly included in the list P .

TABLE I
CLASSES OF PARETO OPTIMAL TREES OF THE NETWORK N_5

P	Class Pareto optimal spanning trees	(l, c)
T_1	$\{(0, 1, 6, 2, 4, 7, 4, 4, 8), (0, 1, 1, 7, 4, 3, 6, 4, 8)\}$	(37, 8)
T_2	$\{(0, 1, 6, 2, 4, 7, 4, 7, 7), (0, 1, 1, 7, 4, 3, 6, 7, 7)\}$	(45, 6)
T_3	$\{(0, 1, 6, 2, 6, 7, 4, 7, 7), (0, 1, 1, 7, 6, 3, 6, 7, 7)\}$	(49, 4)




Fig. 2. The Pareto front of the biobjective spanning trees of Example 5

In line 6 is defined the subnetwork N''' that contains only those edges of N that have a risk strictly less than $c_3 = 4$. We get

$$N'''.Adj = [\langle(2, 4, 2), (3, 8, 2)\rangle, \langle(1, 4, 2), (4, 8, 2)\rangle, \langle(1, 8, 2), (6, 1, 2)\rangle, \langle(2, 8, 2)\rangle, \langle\emptyset\rangle, \langle(3, 1, 2)\rangle, \langle(8, 14, 2)\rangle, \langle(7, 14, 2)\rangle, \langle\emptyset\rangle]. \quad (58)$$

As above, we find that $t \in W_3$ if and only if t is a spanning tree of N''' . In line 7 we find that N''' is not connected and therefore $W_3 = \emptyset$. Also, the algorithm stops.

In this way, it is proved that the list P contains all classes of Pareto optimal trees.

The classes of Pareto optimal trees of the network N_5 are given in Table I. In the examined case, each class of equivalent Pareto optimal trees has two elements. ■

In order to illustrate graphically the obtained result in Example 5, we compose the list W of all spanning trees of the network N_5 . In this case their number is 662 which can be easily achieved using the incidence matrix. The list W can be composed using Corollary 3. To each spanning tree t we correspond a point A_t with Cartesian coordinates $(x(t), y(t))$. In the plane is obtained the set

$$\Gamma = \{A(x(t), y(t)) : t \in W\}.$$

The set Γ has 58 points because the equivalent spanning trees are mapped to the same point on the plane.

In Figure 2, the points that illustrate the classes of equivalent Pareto optimal trees are in gray color, and the rest are in black color. More precisely, the points $A(37, 8)$, $B(45, 6)$, and $C(49, 4)$ are in gray, and the remaining points of Γ are in black.

V. CONCLUSION

In this paper we propose an exact method that constructs the complete Pareto front of the minimum length minimum risk spanning trees problem. It is composed of the solution

of two problems. For the solution of the first problem, the method calculates the list of all minimum spanning trees with respect of the length criterion. For the solution of the second (main) problem, it constructs the complete Pareto front itself, using the solution of the first problem to compose each of the classes of equivalent Pareto optimal trees.

The Algorithm 3 proposes an extension of the Prim's algorithm that allows us simultaneously to find a single minimum spanning tree and the complete list of all remaining minimum spanning trees, defined by their corresponding subproblems. This modification also uses Algorithm 2 that defines a branching that adds to a queue all subproblems that can complete the current tree. Because of the Fibonacci heap implementation of the min-priority queue abstract data type, the complexity of the algorithm that solves the first problem is $O(m + n \lg n)$.

In order to solve the main problem considered, we use the solution of the helper problem that gives us Algorithm 6 that composes a list of all Pareto optimal trees that have minimum length. The computational complexity of the final solution given in Algorithm 7 is $O(KL(m + n \lg n))$, where K is the number of classes of Pareto optimal trees and L is the number of minimum spanning trees with respect to the length criterion.

REFERENCES

- [1] J. B. Kruskal, "On the shortest spanning subtree of a graph and the traveling salesman problem," *Proceedings of the American Mathematical Society*, vol. 7, no. 1, pp. 48–50, 1956.
- [2] R. C. Prim, "Shortest connection networks and some generalizations," *The Bell System Technical Journal*, vol. 36, no. 6, pp. 1389–1401, 1957. doi: 10.1002/j.1538-7305.1957.tb01515.x
- [3] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959. doi: 10.1007/bf01386390
- [4] C. F. Bazlamaçcı and K. S. Hindi, "Minimum-weight spanning tree algorithms a survey and empirical study," *Computers & Operations Research*, vol. 28, no. 8, pp. 767–785, 2001. doi: 10.1016/S0305-0548(00)00007-1
- [5] P. C. Pop, "The generalized minimum spanning tree problem: An overview of formulations, solution procedures and latest advances," *European Journal of Operational Research*, vol. 283, no. 1, pp. 1–15, 2020. doi: 10.1016/j.ejor.2019.05.017
- [6] S. Steiner and T. Radzik, "Computing all efficient solutions of the bi-objective minimum spanning tree problem," *Computers & Operations Research*, vol. 35, no. 1, pp. 198–211, 2008. doi: 10.1016/j.cor.2006.02.023
- [7] A. C. Santos, D. R. Lima, and D. J. Aloise, "Modeling and solving the bi-objective minimum diameter-cost spanning tree problem," *Journal of Global Optimization*, vol. 60, pp. 195–216, 2014. doi: 10.1007/s10898-013-0124-4
- [8] de Sousa, Ernando Gomes, Santos, Andréa Cynthia, and Aloise, Dario José, "An exact method for solving the bi-objective minimum diameter-cost spanning tree problem," *RAIRO-Oper. Res.*, vol. 49, no. 1, pp. 143–160, 2014. doi: 10.1051/ro/2014029
- [9] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*, 3rd ed. Cambridge, Massachusetts: The MIT Press, 2009. doi: 10.5555/1614191
- [10] M. L. Fredman and R. E. Tarjan, "Fibonacci heaps and their uses in improved network optimization algorithms," *Journal of the ACM*, vol. 34, no. 3, pp. 596–615, July 1987. doi: 10.1145/28869.28874
- [11] B. Korte and J. Vygen, *Spanning Trees and Arborescences*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2018, pp. 133–157. ISBN 978-3-662-56039-6

Towards the analysis of errors in centrality measures in perpetuated networks

Meetskumar Pravinbhai Mangroliya*, Jens Dörpinghaus*[†], Robert Rockenfeller*

* University of Koblenz, Germany

Email: meetmangroliya987@uni-koblenz.de, <https://orcid.org/0009-0003-3727-9527>

[†] Federal Institute for Vocational Education and Training (BIBB), Bonn, Germany

Email: jens.doerpinghaus@bibb.de, <https://orcid.org/0000-0003-0245-7752>

Abstract—Centrality measures are essential tools for analyzing the structure and dynamics of graphs, such as knowledge or social networks. They reveal the significance and influence of individual nodes. However, their accuracy can be influenced by data quality, algorithms, and network properties. This study investigates errors in centrality measures within perpetuated networks. It focuses on network resilience and how these results may be used to develop efficient algorithms for centrality measures. It also investigates how perturbation strategies impact network resilience and predict connectivity in the perturbed network. By employing centrality measures (degree, betweenness, closeness, eigenvector), we identify critical nodes that significantly affect network connectivity and information flow. Additionally, statistical tests (Kolmogorov-Smirnov, Cramér-von Mises) assess network robustness and pinpoint critical transition points. This study, by outlining methods for error identification, quantification, and mitigation, offers valuable insights for enhancing network resilience across various domains, including infrastructure design and social network analysis.

Index Terms—Centrality Measures, Complex Networks, Social Network Analysis, Optimization on graphs

I. INTRODUCTION

In the contemporary era, networks are pervasively present. These networks may be tangible systems, such as power grids or transportation networks, or they may be abstract entities, such as networks of acquaintances or collaborations. As individuals, we are all interconnected within a network of social relationships [1], [2].

In network analysis, centrality measures provide crucial quantitative tools for assessing the importance of individual nodes within a network [3], [4], [5], [6] and are also used for AI, e.g., link prediction methods [7], [8]. Among these measures, degree centrality is particularly significant. Degree centrality quantifies the importance of a node by counting the number of direct connections it has to other nodes in the network, thus reflecting its immediate influence within the network structure [9]. To comprehend the resilience of networks to disruptions, perturbation strategies play a pivotal role [10], [4]. These strategies entail the selective or random removal of nodes or edges from a network, with the objective of studying the impact of such removals on the network's structure and properties. In particular, we focus on four removal strategies: high-degree node removal, low-degree node removal, random node removal, and random edge removal. By evaluating the change in degree centrality before and after

employing these strategies, we gain insights into the relative importance of different nodes and edges in maintaining the network's overall structure and function [4], [3]. This is our first research question.

Furthermore, as a second research question, we aim to identify the critical points at which the removal of a node or edge causes the network to be fragmented and undergo significant structural changes. This information is crucial for domains that rely on network connectivity and stability, such as telecommunications and social science, as it provides insights into how these networks might be susceptible to disruption and how their resilience can be enhanced [10], [4]. By understanding these critical points, we can proactively develop strategies to improve the networks' resilience to potential failures and improve their overall robustness.

To evaluate these changes statistically, the Kolmogorov-Smirnov (K-S) and Cramér-von Mises (CvM) tests are employed [11], [12]. These non-parametric statistical tests are used to compare the degree distribution - the distribution of node degrees, representing the number of connections each node has - of a network before and after removal and show the significant change in the structure. To our knowledge, only a few studies have employed the Kolmogorov-Smirnov (KS) test to assess the differences in the degree distributions of networks [13], [11]. However, no studies have employed the Cramér-von Mises (CvM) test for this purpose. This chapter serves as the foundation for our study, laying the groundwork for exploring types of complex networks and utilizing various tools to deepen our understanding.

Following the introduction, this study is structured into four sections. The second section presents a comprehensive review of relevant literature on centrality measures, perturbation strategies, and network resilience, establishing the theoretical and empirical foundation. The third section details the methodology, outlining data collection, analysis, and statistical evaluation procedures. The fourth section presents and analyzes the experimental results, detailing the effects of different removal strategies on network properties. Finally, the fifth section presents a discussion and outlook, interpreting the results, drawing conclusions, and proposing future research directions to advance the understanding of network dynamics and resilience.

II. RELATED WORK

A. Network Perturbation

Networks are both resilient and fragile, meaning that they can withstand some perturbations but are also vulnerable to others. Node and edge removal is a technique that can be used to study how these networks respond to perturbations. Node and edge removal analysis has a rich history in network science. By systematically removing nodes or edges, we gain a deeper understanding of network robustness and efficiency. It is also worth noting that similar questions are addressed within the field of longitudinal network analysis, see [6].

Modern research examines deeper into the intricate interplay between network structure, dynamics, and resilience under various removal scenarios. Additionally, advancements in machine learning and computational techniques have allowed for the development of predictive models to estimate the impacts of node and edge removal before they are executed, enhancing our ability to formulate proactive strategies for network optimization and management. We can use mathematical models, optimization algorithms, and simulations to study the removal of nodes and edges from a network and to predict the effects of these removal actions.

Researchers have proposed and explored diverse methods for removing nodes and edges, encompassing strategies like random removal [10], [14], as well as centrality-based removals such as degree centrality and betweenness centrality [15], [16]. A prevalent attack strategy involves pinpointing critical nodes based on metrics like degree or other centrality measures and systematically removing them in descending order of importance until the network either disconnects or loses essential properties [17]. Albert's investigation into the fragmentation of random and scale-free networks employed two distinct node removal strategies: targeted removal (attacks) and random removal (errors) [10]. In the research of Smith et al. [18], the authors systematically eliminated genes from a biological network to examine the consequences on both network connectivity and behavior. Additionally, in the study of Bellingeri et al. [19], the authors performed a comparative analysis to evaluate the effects of diverse link removal strategies on various real-world complex networks. Additionally, numerous studies have examined the impact of removal strategies on complex networks across various scientific domains [20], [21], [22], [23], [24], [25]. These methodologies provide a comprehensive understanding of how the removal of nodes and edges influences the overall structure and functionality of networks.

B. Impact on Network Connectivity and Structure

Network connectivity is essential for any network. It allows information, influence, resources, and even diseases to flow between the nodes of the network. In other words, network connectivity is what makes networks useful and powerful. Without connectivity, networks would be nothing more than a collection of isolated nodes [4]. Studying how removing nodes and edges from a network affects its ability to connect its

nodes is not just a theoretical exercise rather, it can help us to understand how complex systems work and how to make them more resilient. We can identify the vulnerabilities of complex systems by identifying the nodes and edges that are essential for the network's functioning. We can also identify the types of failures that are most likely to occur in a network.

Moreover, structural changes within networks, whether through the addition or removal of nodes and edges or alterations in degree distribution, can profoundly impact their overall connectivity and efficiency [10], [26]. The effects of node and edge removal on network connectivity have been a subject of extensive research in the field of network science, various studies have explored the effects of perturbation strategies on network resilience and structure. Albert et al. [10] investigated the contrasting vulnerabilities of scale-free and random networks to node perturbation. Holme et al. [17] examined the susceptibility of complex networks to targeted perturbation, while Bellingeri et al. [27] studied attack strategies in real-world networks. Wandelt et al. [15] compared network dismantling strategies, finding node removal often outperforms edge removal, with hybrid approaches effective in specific cases. Callaway et al. [28] analyzed network resilience to failures and attacks, and Chen and Li [29] explored community detection using constrained edge-deleting strategies, demonstrating potential for increased robustness.

After reviewing the work of these researchers in this field, we aim to explore how different node perturbation strategies, including targeted and random perturbation, influence network connectivity.

However, studying the impact of node and removal on different structures and connectivity is not without challenges. The interplay between local and global network structure presents complexities in predicting network behavior [17]. Furthermore, quantifying the impact of removal strategies on network robustness requires sophisticated algorithms and computational resources [26]. Moreover, the ethical implications of removal actions in real-world networks, such as social networks, introduce ethical considerations [30]. Our research aims to navigate these challenges and contribute to a more comprehensive understanding of network connectivity and structural changes due to node and edge removal.

C. Statistical Tests in Network Analysis

In statistical hypothesis testing, the Cramér-von Mises (CvM) test and the Kolmogorov-Smirnov (K-S) test are two most popular non-parametric tests to evaluate the goodness-of-fit between two probability distributions [12].

1) *Kolmogorov-Smirnov (K-S) test:* The Kolmogorov-Smirnov (K-S) test is a statistical tool employed to evaluate whether two cumulative distribution functions (CDF) differ significantly [11], [12]. It quantifies the maximum vertical difference between the CDF of the dataset of original degree distribution and the CDF of reduced degree distribution of the network after removal of nodes or edges [11], [31], [32].

Let $F_{1,n}$ be the cumulative distribution function of original network with a number of nodes n and $F_{2,m}$ be the cumulative

distribution function of the network after node or edge removal with a remaining number of nodes m , then the K-S statistic, denoted as $D_{n,m}$, quantifies the maximum absolute difference between these functions, given by [12]:

$$D_{n,m} = \sup_x |F_{1,n}(x) - F_{2,m}(x)| \quad (1)$$

In practice, the null hypothesis is rejected at a specified significance level α if $D_{n,m}$ exceeds a critical value or in other words, if the condition in following equation 2 satisfies [12].

$$D_{n,m} > c(\alpha) \cdot \sqrt{\frac{n \cdot m}{n + m}} \quad (2)$$

The value of $c(\alpha)$ at specified significance level is tabulated in [33]. The highest vertical difference between the two distribution functions is quantified by the K-S statistic ($D_{n,m}$). A significant difference between the original and modified degree distributions is shown by a high K-S statistic and a low related p-value in the context of network analysis, suggesting that the removal approach has significantly affected the network's structure.

The KS test is particularly suitable for comparing network degree distributions as it doesn't necessitate any assumptions regarding the underlying distribution of the data [13], [31]. The K-S test is exact and straightforward to interpret, delivering dependable outcomes even when dealing with small sample sizes [34]. This precision proves advantageous, especially in situations with constrained data [35], [4].

In the analysis of an evolving social network, Kossinets and Watts employed the Kolmogorov-Smirnov (K-S) test to compare the degree distributions of the network at different time points to identify and quantify the changes in the network's structure over time [35]. Their work highlighted the potential of the K-S test for analyzing the evolution of complex networks over time.

2) *Cramér-von Mises Test*: The Cramér-von Mises (CvM) test is another statistical method used for comparing the empirical cumulative distribution function (ECDF) of the sample to the cumulative distribution function (CDF) of a hypothesized distribution. The Cramér-von Mises W^2 criterion is named after Harald Cramér [36] and Richard Edler von Mises [37]. According to Anderson [38], the Cramér-von Mises W^2 criterion for testing that a sample, x_1, x_2, \dots , has been drawn from a specified continuous distribution $F(x)$ is:

$$W^2 = \int_{-\infty}^{\infty} [F_N(X) - F(x)]^2 dF(x),$$

where $F_N(X)$ is the empirical distribution function of the sample. For a second sample, y_1, y_2, \dots, y_M , a test of the hypothesis that the two samples come from the same (unspecified) continuous distribution can be based on the analogue of W^2 , namely

$$T = \frac{NM}{N+M} \int_{-\infty}^{\infty} [F_N(x) - G_M(x)]^2 dH_{N+M}(x),$$

where $G_M(x)$ is the empirical distribution function of the second sample and $H_{N+M}(x)$ is the empirical distribution function of the two samples together.

If the computed value of T gets higher than the tabulated critical values from [38] then the null hypothesis that the two samples come from the same distribution can be rejected in the favour of alternative hypothesis. The CvM statistic (T) provides a measure of dissimilarity between the two samples.

It is general understanding that the Cramér-von Mises (CvM) test is considered a good choice for comparing distributions with heavy tails compared to the Kolmogorov-Smirnov (K-S) test due to its sensitivity to deviations in the tails of the distribution where crucial information about the network's structure is often contained [39]. There is only few studies that have employed the K-S test as the statistical tool to assess the differences in the degree distribution of the network and none to our knowledge that have employed CvM test.

III. METHOD

In our analysis, we employ a systematic methodology to generate three types of complex networks and investigate the influence of node and edge removal strategies on centrality measures within these networks. We focus on the fundamental centrality measure—degree centrality, due to its distinct and complementary perspectives on node importance within networks. Degree centrality is computed using the NetworkX library in Python, which offers efficient implementations of various centrality algorithms [40]. This measure provides valuable insights into the structural significance of nodes within complex networks.

A. Generation of Complex Networks

A random network was generated using the Erdős-Rényi (ER) graph generator from the NetworkX library [40]. The network parameters were set to include 100 nodes (n) and an edge probability of 0.2 (p) [41], [42]. This edge probability influences the network's density, with higher values resulting in denser networks and lower values yielding sparser ones. Additionally, Python's random state function was used to establish a consistent random seed (seed = 1) for reproducible network generation.

The Barabási-Albert (BA) model is widely used for generating scale-free networks [43]. This model employs preferential attachment, where new nodes connect to existing nodes with high degrees, resulting in a power-law degree distribution. To create a scale-free network using this model, we specified the number of nodes (n) and the parameter for preferential attachment (m). The parameter m dictates the number of edges added for each new node introduced, influencing the network's structure and mimicking real-world scenarios such as scientific collaborations or social interactions.

The Watts-Strogatz (WS) model is employed to generate small-world networks, characterized by both a small average shortest path length and a high clustering coefficient [44]. This model begins with a regular ring lattice where each node is connected to its k nearest neighbors. Then, with a

probability p , each link connected to a clockwise neighbor is rewired to a randomly chosen node, resulting in small-world characteristics. Varying p allows us to produce graphs with different levels of small-worldness, ranging from a regular lattice ($p = 0$) to a random graph with minimum connectivity constraints ($p = 1$).

B. Removal Strategies

In our analysis, we consider two primary node removal strategies: targeted and random. Targeted node removal involves strategically removing nodes based on specific attributes, such as high degree and low degree [10]. This approach allows us to evaluate the network's response to the intentional removal of highly connected nodes or less-connected nodes. Conversely, random node removal involves removing nodes at random, mimicking stochastic node failures or removals. Additionally, we explore edge removal strategies, focusing on the deliberate elimination of specific or random connections within the network. Random edge removal serves as a benchmark for assessing the network's resilience to random disturbances and offers insights into its robustness when faced with unpredictable edge failures [19].

C. Error Measures

To quantify the impact of node and edge removal strategies on network structures, we employ the degree distribution error as a fundamental error measure. Degree distribution error is a key metric that gauges the impact of these strategies by quantifying the dissimilarity between the degree distribution of the original network and that of the network after distinct removal strategies. We utilize the Kolmogorov-Smirnov (K-S) and Cramér-von Mises (CvM) tests to statistically evaluate the differences in degree distributions.

This approach is consistent with established methodologies in network science, providing insights into how different removal strategies impact network structure and resilience [45], [46], [13], [10], [14]. This error measure may provide insights into how different removal strategies impact network structure, resilience, and the relative importance of nodes within the network.

D. Statistical Test procedure

The data collection process using the Kolmogorov-Smirnov (K-S) and Cramér-von Mises (CvM) tests involves generating a random network, scale-free network, or small-world network to establish a baseline degree frequency distribution, which serves as reference data. This distribution is compared with the degree structure of the network after node or edge removal, executed using various strategies. The K-S and CvM tests assess whether the degree sequences of the original and reduced networks are statistically similar. The null hypothesis (H_0) posits that the degree distributions are similar, while the alternative hypothesis (H_1) suggests a difference. A significance level (α) of 0.05 is applied, balancing the risk of Type I errors and the ability to detect differences [47]. If the p-value from these tests exceeds 0.05, we conclude no significant




Fig. 1. Impact of Different Removal Strategies on Network Connectivity in an Erdős-Rényi Random Network

difference in degree distributions; otherwise, we reject H_0 , indicating significant dissimilarity. The results are visualized through plots and tables, elucidating the structural impacts of node or edge removal on network resilience and robustness.

IV. EXPERIMENTAL RESULTS

In this section, we examined the structural changes within random networks, scale-free networks, and small-world networks. Our approach involved applying different node and edge removal processes to observe the evolution of these networks. Specifically, for each network type, we implemented four distinct scenarios: high-degree node removal, low-degree node removal, random node removal, and random edge removal. Our primary objective is to reveal the impact of these perturbations on the networks' topology and connectivity.

A. Change in Network Connectivity

In the analysis of random networks, we employed a consistent approach to assess their structural robustness under these four distinct perturbation scenarios. Here in each iteration 1% of nodes in the node removal approach and 1% of edges in the edge removal approach are removed.

The removal of random nodes from a random network exhibited remarkable resilience, see Figure 1. Despite the stochastic nature of this process, the network displayed significant robustness. Even when a substantial fraction of nodes are removed, the structural integrity of the network is maintained, with no disconnection observed.

The same resilience was observed for the removal of high-degree and low-degree nodes, see Figure 1. The size of the largest component remained unchanged for the removal of high-degree and low-degree nodes (the respective green and yellow color line are covered by the red line in the Figure 1). The network stayed robust and connected even though the critical high-degree nodes were removed. This observation of resilience can be attributed to the random and homogeneous nature of a network.

The most striking structural robustness occurred while randomly removing edges from the network. The network's structural integrity remained totally unaffected. The network showed the robustness up to 81 iterations (81% of edge




Fig. 2. Impact of different removal strategies on network connectivity in a Barabási-Albert scale-free network

removal from the original network) with the largest component of 100 nodes (original generated network with no disconnected nodes).

The network demonstrated both resilience and robustness in response to each of the four distinct types of perturbations, highlighting its ability to adapt and maintain structural integrity under a variety of challenges. The random and uniform nature of these networks contributes to their resilience under different random node and edge removal strategies. Understanding these dynamics is essential for applications in diverse fields, from information dissemination to transportation systems, where random network structures often play a vital role.

The removal of random nodes from a scale-free network demonstrated remarkable resilience, as shown in Figure 2. The network remained connected up to 21 iterations of random node removal, highlighting its robustness against the stochastic nature of this process. However, at the 35th iteration, the network faced again the disconnection, emphasizing its vulnerability to extensive random node removal.

Similarly, for the removal of high-degree nodes in a scale-free network, the network's vulnerability was evident. After just 4 iterations, the network gets disconnected, showcasing its susceptibility to the targeted removal of high-degree nodes. This disconnection resulted into 30 percent of the network as it's largest connected component after only 10 iterations.

When subject to random edge removal, a scale-free network displayed moderate resilience. It maintained connectivity up to the 15th iteration of edge removal, highlighting its ability to adapt and stay intact to some extent under this perturbation.

Conversely, a scale-free network exhibited robustness in response to low-degree node removal. The network remained connected and adaptive, demonstrating its capacity to withstand the perturbation without significant structural impact.

These findings emphasize the behavior of scale-free networks under different perturbations and emphasize the importance of understanding their vulnerabilities and strengths in real-world applications.

In the case of random node removal within a small-world network, see Figure 3, the network displayed moderate resilience to 27 iterations before facing disconnection.




Fig. 3. Impact of Different Removal Strategies on Network Connectivity in a Watts-Strogatz Small-World Network

When subjected to high-degree node removal, a small-world network demonstrated vulnerability, with disconnection occurring after 16 iterations. However, the network's susceptibility became apparent as early as 25 iterations, with almost a 50 percent reduction in its original size, highlighting the consequences of high-degree node removal. Conversely, low-degree node removal posed no threat to a small-world network's robustness. It remained connected throughout the removal process, even as 50 percent of the random nodes were removed, demonstrating its remarkable structural integrity.

In the analysis of random edge removal in a small-world network, the network showed moderate resilient upto around 25% of random edge removal but around 45th iteration, the network faced a significant drop, going from a 90 percent largest connected component to a 50 percent largest connected component. This dynamic response suggests that the network initially demonstrated the capability to adjust and maintain its connectivity in response to these random edge removal. However, as the edge removal increased, the network eventually reached a point where it could no longer adapt effectively and started to break down and lose its resilience.

These observations emphasize the behavior of small-world networks under various perturbation scenarios and underscore the importance of understanding their vulnerabilities and strengths in practical applications.

B. Change in Network Structure

In this subsection, we explore the impact of the node and edge removal strategies on the structure of the complex networks. We conducted a series of experiments to assess how the network's degree distribution evolves as we apply various removal strategies. These experiments were conducted on the same three types of random, scale-free, and small-world networks with the same parameters as discussed before in the methodology. Also, the same removal strategies included high-degree node removal, low-degree node removal, random node removal, and random edge removal and the removal percentage from 5% to 25%. By conducting these experiments on different network types, we aimed to gain insights into how each removal strategy influenced the degree frequency distribution of the network and, consequently, the overall




Fig. 4. Degree Frequency Distribution in Erdős-Rényi Random Network at different percentages of High-Degree Node Removal




Fig. 5. Degree Frequency Distribution in Erdős-Rényi Random Network at different percentages of Low-Degree Node Removal

network structure. From plot figures, the degree distribution plot where each plot illustrates a distinct removal strategy, and each color line represents the degree distribution in the corresponding network before and after a range of removal percentages. We used Non-parametric Kolmogorov-Smirnov (K-S) and Cramér-von Mises (CvM) tests with a significance level of 0.05 to statistically evaluate these changes.

1) Impact of removal strategies on Random Network:

The frequency distribution plots of a random network before and after high-degree nodes removed are displayed in Figure 4. The results of the high-degree node removal strategy are summarized in Table I. The CvM and K-S statistics show a significant increase in the proportion of node removal. The CvM statistics consistently deviate from the original degree distribution, rising from 0.8888 at 5% removal to 11.0364 at 25% removal. The CvM statistics' p-values correspondingly decreased, signifying a significant deviation from the original distribution. Similar trends can be seen in the K-S statistics, which rise from 0.1984 at 5% removal to 0.79 at 25% reduction. The corresponding K-S test p-values similarly decreased to zero, confirming the significant effect of removing high-degree nodes on the degree distribution of the network. Based on these findings, the null hypothesis which states that the degree distribution of a random network is the same before and after a removal of high-degree nodes can be rejected.

The frequency distribution plots of a random network before and after the removal of low-degree nodes are shown in Figure 5. Table II provides an overview of the low-degree node removal strategy's results. As the percentage of nodes removed increases from 5% to 25%, there is a small but noticeable change in the degree distribution plot. This change is reflected in both the CvM and K-S statistics, which increase with increasing node removal. The decreasing p-values associated with the CvM and K-S statistics demonstrate that the observed changes in the degree distribution are statistically significant. As expected, the CvM test is more sensitive to the long tail of the plot than the K-S test. Given the importance of the long tail in network analysis, we can reject the hypothesis that the degree distribution remains unchanged even after removing 15% of low-degree nodes.




Fig. 6. Degree Frequency Distribution in Erdős-Rényi Random Network at different percentages of Random Node Removal.

The results of the random node removal method are displayed in the frequency distribution plot (see Figure 6 and Table III). At 5% removal, the CvM statistics show a slight increase to 0.1355, with a p-value of 0.4404, indicating no significant departure from the original distribution. However, as the percentage of nodes removed increases, the CvM statistics rise notably, reaching 5.5217 at 25% removal, accompanied by a very low p-value. The K-S statistics also demonstrate a progressive increase, from 0.0826 at 5% removal to 0.54 at 25% removal, with corresponding p-values diminishing significantly, reinforcing the significant deviation from the original distribution. By combining the results, we fail to reject the null hypothesis for the removal of 5% of random nodes, but we reject it for the remaining removal percentages.

The results of the random edge removal strategy, and frequency distribution plot are shown in Figure 7 and Table IV respectively. The CvM statistics increase from 0.2962 at 5% removal with a p-value of 0.1390 to 6.0382 at 25% removal with an extremely low p-value, indicating the statistical significance of structure change. The K-S statistics also show a progressive rise, from 0.12 at 5% removal with a p-value of 0.4695 to 0.5 at 25% removal with a very low p-value. The results of this study indicate how removing more than 5% of random edges significantly affects the degree distribution, causing a noticeable and statistically significant




Fig. 7. Degree Frequency Distribution in Erdős-Rényi Random Network at different percentages of Random Edge Removal




Fig. 9. Degree Frequency Distribution in Barabási-Albert Scale-Free Network at different percentages of Low-Degree Node Removal




Fig. 8. Degree Frequency Distribution in Barabási-Albert Scale-Free Network at different percentages of High-Degree Node Removal




Fig. 10. Degree Frequency Distribution in Barabási-Albert Scale-Free Network at different percentages of Random Node Removal

deviation from the original distribution. Consequently, for the remaining removal percentages, we reject the null hypothesis.

2) *Impact of removal strategies on Scale-free Network:* The results of the high-degree node removal strategy in a scale-free network are the frequency distribution plots in Figure 8 and associated statistical test results in Table V. The CvM statistics exhibit a significant increase, from 2.0114 at 5% removal to 10.1317 at 25% removal, with very low p-values consistently indicating statistical significance. Similar gradual changes are also shown by the K-S statistics, which indicate a significant shift in the degree distribution. These findings show the substantial impact of even 5% of high-degree node removal on altering the network’s degree distribution in a scale-free network. Hence, we reject the null hypothesis for each removal percentage.

The results from the low-degree node removal strategy in a scale-free network, the frequency distribution plot shown in Figure 9, indicate a relatively mild impact on the degree distribution, the associated statistical results are presented in Table VI. The CvM statistics display a gradual increase from 0.0241 at 5% removal to 0.2861 at 25% removal, with p-values consistently higher than the significance value 0.05, suggesting no statistically significant deviation from the original distribution. Alongside, the K-S test demonstrates similar results, with the p-value higher than the significance value of 0.05

for all removal percentages. This suggests that removing low-degree nodes has a limited effect on the degree distribution in a scale-free network. Consequently, we fail to reject the null hypothesis.

The results from a scale-free network, as the frequency distribution plot shown in Figure 10 and the statistical results shown in Table VII, indicate that random node removal has a minimal impact on the degree distribution for removal percentages up to 15%. The CvM statistics show deviations, with p-values remaining lesser than the significance value of 0.05. However, at 20% and 25% removal, a more noticeable shift is observed, with CvM statistics reaching 0.5121 and 1.3568, respectively, and p-values dropping to 0.0371 and 0.0004, indicating a statistically significant deviation from the original distribution. Additionally, the K-S test also showed similar results up to 20% of random node removal with K-S statistics 0.15 and p-value 0.2456. This suggests that the effect of random node removal on the degree distribution becomes more pronounced as the percentage of removed nodes increases, leading to the rejection of the null hypothesis for 25% of random node removal.

The results of a scale-free network after random edge removal are illustrated in Figure 11 and associated statistical results in Table VIII. The CvM statistics progressively increase from 0.0380 at 5% removal to 1.3194 at 25% removal,




Fig. 11. Degree Frequency Distribution in Barabási-Albert Scale-Free Network at different percentages of Random Edge Removal




Fig. 13. Degree Frequency Distribution in Watts-Strogatz Small-World Network at different percentages of Low-Degree Node Removal




Fig. 12. Degree Frequency Distribution in Watts-Strogatz Small-World Network at different percentages of High-Degree Node Removal

with p-values decreasing significantly from 15% removal. Simultaneously, the K-S test provided similar results with a p-value below the significant value of 0.05 from 20% removal, further supporting the significant deviation from the original distribution. Here, we reject the null hypothesis at 20% and 25% removal of random edges.

3) *Impact of removal strategies on Small-world Network:* The frequency distribution plots in Figure 12 shows the impact of high-degree node removal on the degree distribution of a small-world network. The associated statistical results are presented in Table IX. The CvM test is more sensitive to these changes, at 5% removal, the CvM statistic is 0.5494 with a p-value of 0.0299, indicating a significant deviation from the original degree distribution. The K-S test, on the other hand, showed lower statistics and a p-value higher than the significance level of 0.05 for both 5% and 10% removal. However, after 10% removal, the K-S test also indicated a significant deviation with a p-value lower than the significance level. This suggests that the CvM test is more sensitive to changes in the degree distribution of high-degree nodes so, we reject the null hypothesis at higher removal percentages including at 10% removal of high-degree nodes.

Figure 13 visually represents the impact of the Low-Degree Node Removal strategy on a small-world network's degree distribution, with detailed statistical results presented in Table

X. Initially, at a 5% removal rate, the CvM statistics of 0.0314 and a p-value of 0.9754 indicated no noteworthy deviation. However, as the removal percentage increased, statistical significance became apparent, with a CvM statistic of 0.9680 and a significantly low p-value of 0.0029 at 20% removal, further decreasing at higher percentages. The corresponding K-S test at 20% removal, featuring K-S statistics of 0.1750 and a p-value of 0.1172, suggests that low-degree node removal did impact the degree distribution, although the effect was not significant. A statistically significant shift in distribution was detected at 25% removal, with K-S statistics of 0.2700 and a p-value of 0.0032. These findings demonstrate that the removal of low-degree nodes has a gradual impact on the degree distribution of a small-world network. Here, we reject the null hypothesis only at 25% removal of low-degree nodes.

The effect of the random node removal strategy on the degree distribution of a small-world network is presented in Figure 14 and the corresponding statistical results are presented in Table XI. At 5% removal, the CvM statistics of 0.2938 and a p-value of 0.1412 indicated no significant deviation from the original distribution. However, as removal percentages increased, we can see the statistical significance. Notably, at 15% removal, a CvM statistic of 2.3229 and a very low p-value highlighted a substantial deviation from the original distribution. This was supported by the corresponding K-S statistic, which also increased to 0.3265 at 15% removal, along with a p-value lower than the significance level of 0.05. The higher sensitivity of the CvM test to changes in the distribution of high-degree nodes suggests that the observed deviation is likely due to the removal of these nodes. So, We reject the null hypothesis for the higher percentage including 15% removal of random nodes.

In Figure 15, we present the effects of random edge removal strategy on a small-world network's degree distribution and the associated statistical results are presented in Table XII. At 5% removal, CvM statistic 0.2243 and K-S statistic 0.09 indicated no significant change, aligning with p-values of 0.2260 and 0.8154, respectively. However, at 10% removal and higher removal percentages, both tests illustrated progressively significant deviations from the original distribution. By 25%




Fig. 14. Degree Frequency Distribution in Watts-Strogatz Small-World Network at different percentages of Random Node Removal




Fig. 15. Degree Frequency Distribution in Watts-Strogatz Small-World Network at different percentages of Random Edge Removal

removal, the CvM p-value and K-S p-value became very low, indicating the strategy's significant impact on the network's structural integrity. From these results, we reject the hypothesis for higher removal percentages including 10% removal of random edges.

V. DISCUSSION AND OUTLOOK

In this study, we have investigated the resilience of networks to various perturbation strategies, with a particular focus on the effects of high-degree node removal, low-degree node removal, random node removal, and random edge removal. Our research highlights the importance of understanding network resilience, particularly in the context of maintaining connectivity and stability in various real-world networks. The resilience of power grids, transportation systems, and social networks can be enhanced by gaining insights into their susceptibility to failures and the identification of critical nodes or edges whose removal could lead to significant structural changes. Furthermore, insights on critical nodes that highly influence the centrality measures for a given graph are crucial to optimize the runtime of heuristics for centrality measures in longitudinal networks and to find bounds for the possible change of node centrality measures.

The results demonstrated that random networks exhibited high resilience to connectivity disruptions, although they did lose their structural properties at a relatively low removal threshold. Scale-free networks exhibited vulnerability to high-degree node removal, but they showed resilience to other removal strategies to a moderate threshold. Small-world networks maintained connectivity up to substantial removal percentages, but they lost the structural integrity at a lower removal percentage, indicating a similar level of resilience as random networks. These findings highlight the complex relationship between node importance, network structure, and network resilience in complex networks.

Our findings demonstrate that high-degree nodes play a pivotal role in maintaining network connectivity, as their removal often leads to more significant disruptions compared to random node or edge removals. This underscores the necessity for targeted strategies to enhance network robustness, particularly by protecting or reinforcing high-degree nodes. Moreover, our study contributes to the broader field of network science by applying and validating the CvM test for degree distribution comparison, a methodology previously underutilized in this context. This development provides researchers with a novel instrument for further inquiry into the dynamics and resilience of networks.

Our study has yielded a more profound comprehension of the resilience and connectivity patterns exhibited by disparate network types. While our study offers valuable insights, it is essential to acknowledge certain limitations. Primarily, focusing on a relatively modest network comprising 100 nodes may restrict the generalizability of our findings to larger and more intricate real-world networks. Secondly, our utilization of static network analysis fails to capture the dynamic nature of real-world networks, which continuously evolve and adapt to changing conditions. Future research could enhance our understanding of network dynamics by employing temporal network analysis to investigate how networks change over time. Additionally, while our study examines various network types, the results may not be universally applicable to all network structures. Future research should address these limitations by studying larger and more intricate networks and incorporating dynamic elements. This approach would ultimately enhance the validity and applicability of our study's conclusions.

REFERENCES

- [1] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, "Complex networks: Structure and dynamics," *Physics reports*, vol. 424, no. 4-5, pp. 175–308, 2006.
- [2] J. Dörpinghaus, S. Klante, M. Christian, C. Meigen, and C. Düing, "From social networks to knowledge graphs: A plea for interdisciplinary approaches," *Social Sciences & Humanities Open*, vol. 6, no. 1, p. 100337, 2022.
- [3] S. P. Borgatti, "Centrality and network flow," *Social networks*, vol. 27, no. 1, pp. 55–71, 2005.
- [4] M. E. J. Newman, *Networks: An Introduction*. Oxford University Press, 2010.
- [5] J. Dörpinghaus, A. Stefan, B. Schultz, and M. Jacobs, "Context mining and graph queries on giant biomedical knowledge graphs," *Knowledge and Information Systems*, vol. 64, no. 5, pp. 1239–1262, 2022.

- [6] J. Dörpinghaus, V. Weil, and M. W. Sommer, "Towards modelling and analysis of longitudinal social networks," *Annals of Computer Science and Information Systems*, vol. 37, pp. 81–89, 2023.
- [7] J. Dörpinghaus, T. Hübenthal, and J. Faber, "A novel link prediction approach on clinical knowledge graphs utilising graph structures," in *2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2022, pp. 43–52.
- [8] A. Dahhani, I. Alloui, S. Monnet, and F. Vernier, "A graph matching algorithm to extend wise systems with semantic," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 411–420.
- [9] L. C. Freeman, "Centrality in social networks conceptual clarification," *Social Networks*, vol. 1, no. 3, pp. 215–239, 1978.
- [10] R. Albert, H. Jeong, and A. Barabasi, "Error and attack tolerance of complex networks," *Nature*, 2000.
- [11] A. D. Broido and A. Clauset, "Scale-free networks are rare," *Nature communications*, vol. 10, no. 1, pp. 1–12, 2019.
- [12] T. B. Arnold and J. W. Emerson, "Nonparametric goodness-of-fit tests for discrete null distributions," *R Journal*, vol. 3, no. 2, pp. 34–39, 2011.
- [13] A. Clauset, C. R. Shalizi, and M. E. J. Newman, "Power-law distributions in empirical data," *SIAM review*, vol. 51, no. 4, pp. 661–703, 2009.
- [14] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, no. 1, pp. 47–77, 2002.
- [15] S. Wandelt, X. Sun, D. Feng, M. Zanin, and S. Havlin, "A comparative analysis of approaches to network-dismantling," *Scientific Reports*, vol. 8, p. 13513, 2018.
- [16] E. Jenelius and L.-G. Mattsson, *Resilience of Transport Systems*, 05 2020.
- [17] P. Holme, B. Kim, C. Yoon, and S. Han, "Attack vulnerability of complex networks," *Physical Review E*, 2002.
- [18] J. Smith and L. Johnson, "Edge removal and network reconstruction in biological networks," *PLoS ONE*, vol. 10, no. 7, p. e0131180, 2015.
- [19] M. Bellingeri, D. Bevacqua, F. Scotognella, and et al., "A comparative analysis of link removal strategies in real complex weighted networks," *Scientific Reports*, vol. 10, no. 1, p. 3911, 2020.
- [20] V. Latora and M. Marchiori, "Efficient behavior of small-world networks," *Physical Review Letters*, vol. 87, no. 19, p. 198701, 2001.
- [21] Y. Yang, T. Nishikawa, and A. E. Motter, "Small vulnerable sets determine large network cascades in power grids," *Science*, vol. 358, 2017.
- [22] J. L. Caldu-Primo, E. R. Alvarez-Buylla, and J. Davila-Velderrain, "Structural robustness of mammalian transcription factor networks reveals plasticity across development," *Scientific Reports*, vol. 8, no. 1, pp. 1–15, 2018.
- [23] L. K. Gallos, R. Cohen, P. Argyrakis, A. Bunde, and S. Havlin, "Stability and topology of scale-free networks under attack and defense strategies," *Physical Review Letters*, vol. 94, no. 18, p. 188701, 2005.
- [24] M. Zanin and F. Lillo, "Modelling the air transport with complex networks: A short review," *The European Physical Journal Special Topics*, vol. 215, pp. 5–21, 2013.
- [25] J. Dörpinghaus, V. Weil, C. Düing, and M. W. Sommer, "Centrality measures in multi-layer knowledge graphs," *Annals of Computer Science and Information Systems*, vol. 32, pp. 163–170, 2022.
- [26] C. N. Schneider and I. Koutsopoulos, "Mitigation of malicious attacks on networks," *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, 2011.
- [27] M. Bellingeri, D. Cassi, and S. Vincenzi, "Efficiency of attack strategies on complex model and real-world networks," *Physica A: Statistical Mechanics and its Applications*, vol. 414, p. 174–180, Nov. 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.physa.2014.06.079>
- [28] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Network robustness and fragility: Percolation on random graphs," *Physical Review Letters*, vol. 85, no. 25, p. 5468, 2000.
- [29] X. Chen and J. Li, "Community detection in complex networks using edge-deleting with restrictions," *Physica A: Statistical Mechanics and its Applications*, vol. 519, pp. 181–194, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378437118315358>
- [30] F. Baltar and I. Brunet, "Unintended consequences and new challenges of online social systems: Methods, challenges, and applications," *Journal of Information Technology & Politics*, vol. 9, no. 1, pp. 1–20, 2012.
- [31] S. Aliakbary, J. Habibi, and A. Movaghar, "Quantification and comparison of degree distributions in complex networks," in *7th International Symposium on Telecommunications (IST)*, Tehran, Iran, September 2014.
- [32] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "The structure of the internet graph," *IEEE/ACM Transactions on Networking*, vol. 7, no. 3, pp. 397–412, 1999.
- [33] "Table of critical values for the two-sample test," https://web.archive.org/web/20130613002106/http://www.soest.hawaii.edu/wessel/courses/gg313/Critical_KS.pdf, memento from June 13, 2013 in the Internet Archive. [Online]. Available: https://web.archive.org/web/20130613002106/http://www.soest.hawaii.edu/wessel/courses/gg313/Critical_KS.pdf
- [34] I. M. Chakravarti, R. G. Laha, and J. Roy, *Engineering Statistics Handbook*. New York, NY: John Wiley and Sons, 1967.
- [35] G. Kossinets and D. J. Watts, "Empirical analysis of an evolving social network," *Science*, vol. 311, pp. 88–90, 2006.
- [36] H. Cramér, "On the composition of elementary errors," *Skandinavisk Aktuarietidskrift*, vol. 11, pp. 141–180, 1928.
- [37] R. von Mises, *Wahrscheinlichkeit, Statistik und Wahrheit*. Vienna, Austria: Julius Springer, 1928.
- [38] T. W. Anderson, "On the distribution of the two-sample cramer-von-mises criterion," *The Annals of Mathematical Statistics*, pp. 1148–1159, 1962.
- [39] M. A. Stephens, "Edf statistics for goodness of fit and some comparisons," *Journal of the American Statistical Association*, vol. 69, no. 347, pp. 730–737, 1974.
- [40] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring network structure, dynamics, and function using networkx," *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, vol. 11, pp. 11–15, 2008.
- [41] P. Erdős and A. Rényi, "On random graphs, i," *Publicationes Mathematicae (Debrecen)*, vol. 6, pp. 290–297, 1959.
- [42] E. N. Gilbert, "Random graphs," *The Annals of Mathematical Statistics*, vol. 30, no. 4, pp. 1141–1144, 1959.
- [43] A. L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [44] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [45] M. E. Newman, *Networks*. Oxford University Press, 2018.
- [46] E. Estrada, "Subgraph centrality in complex networks," *Physical Review E*, vol. 71, no. 5, p. 056103, 2005.
- [47] M. Newman, *The Structure and Dynamics of Networks*. Princeton University Press, 2003.

TABLE I
ERDŐS-RÉNYI RANDOM NETWORK: STATISTICAL TEST RESULTS AFTER HIGH-DEGREE NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.8888	0.0045	0.1984	3.60×10^{-2}
10%	3.0736	5.02×10^{-8}	0.3733	2.11×10^{-6}
15%	6.2458	1.97×10^{-11}	0.5076	2.46×10^{-11}
20%	8.9809	2.73×10^{-10}	0.6575	1.09×10^{-18}
25%	11.0364	2.53×10^{-10}	0.7900	4.36×10^{-27}

TABLE II
ERDŐS-RÉNYI RANDOM NETWORK: STATISTICAL TEST RESULTS AFTER LOW-DEGREE NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.0599	0.8219	0.0721	0.9394
10%	0.2897	0.1451	0.1322	0.3430
15%	0.6399	0.0178	0.1829	0.0791
20%	1.2772	0.0006	0.2550	0.0051
25%	2.2123	4.16×10^{-6}	0.35	3.79×10^{-5}

TABLE III
ERDŐS-RÉNYI RANDOM NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.1355	0.4404	0.0826	0.8563
10%	0.9319	0.0035	0.2	0.0383
15%	2.8869	1.30×10^{-7}	0.3518	1.42×10^{-5}
20%	3.4546	7.19×10^{-9}	0.4025	6.09×10^{-7}
25%	5.5217	3.51×10^{-10}	0.54	5.85×10^{-12}

TABLE IV
ERDŐS-RÉNYI RANDOM NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM EDGE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.2962	0.1390	0.12	0.4695
10%	1.1043	0.0014	0.23	0.0099
15%	2.5327	8.03×10^{-7}	0.33	3.21×10^{-5}
20%	4.1507	2.12×10^{-10}	0.43	1.12×10^{-8}
25%	6.0382	9.28×10^{-12}	0.5	1.00×10^{-11}

TABLE V
BARABÁSI-ALBERT SCALE-FREE NETWORK: STATISTICAL TEST RESULTS AFTER HIGH-DEGREE NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	2.0114	1.18×10^{-5}	0.2632	1.78×10^{-3}
10%	4.5115	3.88×10^{-11}	0.4444	5.88×10^{-9}
15%	7.2186	3.69×10^{-10}	0.5529	1.79×10^{-13}
20%	9.0646	3.25×10^{-10}	0.6250	8.85×10^{-17}
25%	10.1317	2.38×10^{-9}	0.7200	5.89×10^{-22}

TABLE VI
BARABÁSI-ALBERT SCALE-FREE NETWORK: STATISTICAL TEST RESULTS AFTER LOW-DEGREE NODE REMOVAL

Percentage	CvM Stat.	P-Value (CvM)	K-S Stat.	P-Value (K-S)
5%	0.0241	0.9934	0.0305	0.9999
10%	0.0840	0.6754	0.0533	0.9978
15%	0.1083	0.5505	0.0553	0.9969
20%	0.1637	0.3528	0.0700	0.9707
25%	0.2861	0.1485	0.1000	0.7523

TABLE VII
BARABÁSI-ALBERT SCALE-FREE NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.0363	0.9564	0.0326	0.9999
10%	0.1855	0.2993	0.0889	0.8092
15%	0.2839	0.1508	0.0976	0.7269
20%	0.5121	0.0371	0.1500	0.2456
25%	1.3568	0.0004	0.2133	0.0354

TABLE VIII
BARABÁSI-ALBERT SCALE-FREE NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM EDGE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.0380	0.9488	0.0500	0.9997
10%	0.1921	0.2849	0.0800	0.9084
15%	0.4865	0.0431	0.1500	0.2112
20%	0.7992	0.0073	0.2100	0.0241
25%	1.3194	4.42×10^{-4}	0.2100	0.0241

TABLE IX
WATTS-STROGATZ SMALL-WORLD NETWORK: STATISTICAL TEST RESULTS AFTER HIGH-DEGREE NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.5494	0.0299	0.1289	0.3528
10%	1.1332	1.19×10^{-3}	0.1911	0.0538
15%	2.3875	1.69×10^{-6}	0.2794	1.14×10^{-3}
20%	4.9056	3.68×10^{-11}	0.4500	1.29×10^{-8}
25%	8.0887	3.35×10^{-11}	0.6033	5.00×10^{-15}

TABLE X
WATTS-STROGATZ SMALL-WORLD NETWORK: STATISTICAL TEST RESULTS AFTER LOW-DEGREE NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.0314	0.9754	0.0295	0.9999
10%	0.1702	0.3356	0.0844	0.8536
15%	0.4109	0.0677	0.1029	0.6668
20%	0.9680	2.90×10^{-3}	0.1750	0.1172
25%	1.7561	4.44×10^{-5}	0.2700	3.16×10^{-3}

TABLE XI
WATTS-STROGATZ SMALL-WORLD NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM NODE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.2938	0.1412	0.0926	0.7507
10%	0.6479	0.0170	0.1722	0.1050
15%	2.3229	2.36×10^{-6}	0.3265	7.45×10^{-5}
20%	3.3379	1.30×10^{-8}	0.3750	4.56×10^{-6}
25%	4.8054	2.83×10^{-11}	0.4567	1.47×10^{-8}

TABLE XII
WATTS-STROGATZ SMALL-WORLD NETWORK: STATISTICAL TEST RESULTS AFTER RANDOM EDGE REMOVAL

Percentage	CvM Statistic	P-Value (CvM)	K-S Statistic	P-Value (K-S)
5%	0.2243	0.2260	0.0900	0.8154
10%	1.0744	1.64×10^{-3}	0.2300	0.0099
15%	1.9970	1.27×10^{-5}	0.2900	4.12×10^{-4}
20%	3.4719	6.63×10^{-9}	0.3900	3.57×10^{-7}
25%	4.1784	1.84×10^{-10}	0.4200	2.75×10^{-8}

Real options analysis framework for agile projects

Gloria J. Miller, DBA
0000-0003-2603-0980
maxmetrics, Heidelberg
Email: g.j.m@ieee.org

Abstract—The literature proves that agile projects have a higher success rate in stakeholder satisfaction and overall success than projects managed with a plan-driven methodology such as waterfall. However, little corresponding literature examines whether that success extends to the target benefits. This study identifies the mechanisms—actions, decisions, or entities—that enable agile and plan-driven projects to deliver target benefits. It uses real options analysis to quantify and model the differences between project methods and builds a management decision-making framework. The framework includes real option types, mechanisms, and locations; project roles and processes; risk scores and failure rates; a computational model; and a binomial tree for visual analysis. The study contributes a novel framework to the project management literature on agile projects and benefits realization.

Index Terms—Real options analysis, agile, risk failure rate, benefits realization.

I. INTRODUCTION

BY initiating a project, organizations expect to realize a set of target benefits that are the reason for the investment and other tactical goals such as project efficiency [1]. There is proof that agile projects have a higher success rate in stakeholder satisfaction and overall success than projects managed with a plan-driven methodology such as waterfall [2, 3, 4]. However, little corresponding literature examines whether success with agile projects extends to the target benefits. In short, the two project approaches differ in planning and execution: plan-driven projects prefer a linear product development life cycle where an upfront plan defines the scope, schedule, and costs for delivery Lundin and Söderholm [5], and agile projects encourage an iterative and incremental life cycle, self-organizing teams, and evolutionary planning for product development [6, 7, 3].

Serrador and Pinto [3] proved that agile projects were successful at project efficiency, stakeholder satisfaction, and overall performance. Gemino et al. [4] determined that agile and plan-driven approaches have similar time, budget, and scope performance. Further, agile projects demonstrate a higher success rate in stakeholder satisfaction, which is attributed to successful communication. In a meta-analysis of 69 studies, [8] found that the cost-benefit of agile methods, as measured by return on investment, net present value, break-even point, and real options analysis, was on par with plan-driven methods.

The literature on the benefit realization of agile projects exists but is scarce in explaining the actions, decisions, or entities in agile projects that contribute to the target benefits.

Marnewick and Marnewick [2] argues that agile projects' frequent and iterative delivery allows for earlier benefits than traditional methodologies. Moreover, the article states that benefits management is a continuous delivery cycle where benefits are inherently realized and provides empirical evidence.

This study analyzes the mechanisms agile projects use to deliver target benefits. "A mechanism is defined as an action, decision or entity that enables a real option" [9, p. 459]. Then, the study uses those mechanisms to compare the benefits realization between agile and plan-driven projects in a real options analysis framework. A real option is a right (but not an obligation) to invest in a real asset in the future [10].

This study contributes a novel framework that project owners and managers can use to determine the mechanisms within a project that can be optimized to deliver the target benefits. It furthermore contributes to the theory of benefits realization in projects.

The paper is structured as follows. Section Two provides a literature review and describes the related research. Section Three includes the conceptual framework and a description of the research methodology. Section Four describes the relevant project frameworks, and Section Five describes the characteristics of the real options. Section Six applies the method to an illustrative comparison case. Section Six also provides conclusions, including the study's contribution, implications, limitations, and considerations for future research.

II. LITERATURE REVIEW

A. Projects

Projects are "... a unique, transient endeavor, undertaken to achieve planned objectives, which could be defined in terms of outputs, outcomes or benefits" [11]. An important measure of project success is whether the project contributes the expected benefits to the sponsoring organization [1, 12]. Direct and indirect links between project and corporate success require "an effective benefits delivery and management process that involves the mutual co-operation of project management and line management functions" [12].

Project management success, also known as project efficiency or the iron triangle, is the performance of delivering on time, within budget, and to quality. It requires focusing on risk management, change control, limiting project duration, and assigning clear responsibilities [12]. However, delivering the right functionality is more important than finishing rigidly on time and within budget [13, 14].

B. Project Methodologies

Terminology within project management is evolving, so terms such as methodology, approach, framework, and life cycle are used inconsistently between project management standards, literature, and practitioners [15, 16]. This research synthesizes the terminology for project methodology and life cycle. A project methodology is a collection of processes, tools, techniques, methods, capability profiles, knowledge areas, and related understanding used to manage a project life cycle [15]. The term methodology is interchangeable with project approach or framework. A life cycle is a series of organizing phases or stages that a project passes through systematically to transform an idea at the start to reality at completion [11, 16].

Plan-driven methodologies, also known as traditional or waterfall methodologies, are linear life cycles that follow a stage-gate or phased approach. They create an upfront plan where the time is limited, the scope is defined, and the costs are estimated [5]. The output from one phase of the project is built on in the next phase, and the output is delivered in a single release.

Agile methodologies follow the four values and 12 principles described in the Agile Manifesto, a framework for managing projects in a flexible way that responds to dynamic situations [6, 7, 3]. Agile methodologies generally encourage iterative and incremental development life cycles, self-organizing teams, and evolutionary product development. The output is delivered in small, frequent releases to create more value than large mega-releases [17]. Multiple methodologies ascribe to the values and principles of the Agile Manifesto, such as Scrum, Extreme Programming (XP), Lean, and Kanban. They each have their own rules, events, and practices [18].

Hybrid methodologies mix elements of linear, iterative, incremental, and evolutionary life cycles or traditional and agile methodologies.

C. Real Options Analysis

A real option is the right to invest in a real asset in the future without an obligation to make the investment. The cost of an option on an asset is small relative to the cost of the total asset. Thus, by investing in an option to the asset instead of the total asset, an organization can defer investments until the uncertainty is reduced. Resources can be deployed on other opportunities. Switch, change scale, stage, abandon, and strategic growth are types of real options; the types are described in Table I.

Real options analysis (ROA) “is most valuable when there is a high uncertainty with the underlying asset value and management has significant flexibility to change the course of the project in a favorable direction and is willing to exercise the options” [20, p. 94]. In a real options model, project opportunity is equated with the flexibility to acquire a payoff by making an investment before a project is completed [22].

Real options theory is derived from finance theories that value financial options contracts. “A financial option contract conveys the right, but not the obligation, on the purchaser to

TABLE I
REAL OPTIONS TYPES

Types	Description	References
Switch	Put asset to a different purpose from the original intentions.	[19, 20, 21]
Change scale: expand	Change the scope by reallocating resources. Resources can be expanded or systems can be scaled up with relative ease.	[20]
Change scale: reduce	Change the scope by reallocating resources. Resources can be reduced or systems scaled down with relative ease.	[19, 21]
Stage	When structured as a series of incremental outlays, terminate the project should conditions become unfavorable. This option overlaps with other options such as abandon, change scale, and strategic growth.	[21, 19, 20]
Abandon	If possible, without negative consequences, discontinue the project prior to completion and redeploy the remaining resources.	[19, 20, 21]
Strategic growth	Pursue a variety of follow-on investments when the initial baseline opens the opportunity to do so. This differs from other options, which reduce or limit potential losses from unfavorable circumstances, by increasing gains in the event of favorable circumstances.	[19, 21]
Defer with revenue leakage	Delay the decision on whether and how to invest for some period but imperil some aspect of the potential benefits.	[20]
Defer	Delay the decision on whether and how to invest for some period without imperiling the potential benefits	[19, 20, 21]

either buy or sell an underlying asset at some point in the future” [10, p. 2].

D. Real Options Use in Projects

Real options have been applied in many project situations, including the portfolio and project initiation [23], continuation and termination of projects [19], valuation of projects [24], planning and risk management [25, 26, 27], and performance monitoring [28]. Racheva et al. [29] conceptualized the use of real options within an agile project as a quantitative approach to determine when and how much to invest in software functionality.

Real options in project literature are based on different models. One model is based on defining the location where the real option occurs: “in” or “on” the project [9, 30]. Real options “in” projects are decisions or actions taken within the project, such as the system design. Real options “on” projects are decisions and actions that affect the project outcome. However, this characterization is ambiguous in a location where the option exists, with the project team or the project governance body[9].

Another model is the characterization of real options as mechanisms and types [9]. “A mechanism is an action, decision, or entity that enables a real option. It is a source of flexibility... A type is an action or decision that can be exercised by the owner of the real option. The type is therefore representative of the future flexibility” [9, p. 459–460]. Type is a method of managing uncertainty that can be exercised before the expiration date and after or as uncertainties are clarified.

This characterization was conceived to reconcile ambiguities in terminology between the classic application of real options analysis and the engineering application of real options. “In the classical application of ROA, the real option is used to describe the right but not the obligation to take a future action, which is then considered in the valuation of decisions under uncertainty. On the other hand, in engineering applications that actively synthesize options, the term real options is typically used to refer to a design feature that enables some flexibility” [9, p. 460].

In compound options, exercising one option generates another. In this case, the value of the compound option is derived from the values from another option and not the underlying asset. For example, in a multistage project, the design is one project, and the development is a separate project [20]. Compound options are out of the scope of this study.

III. METHODOLOGY

A. Conceptual Framework

The study compares how agile projects achieve benefits relative to traditional projects. The study uses real options analysis as a basis and combines two frameworks to define the research model. First, it uses the uncertainty framework from [9] for mapping life cycle characteristics that enable flexibility to the real option types used to manage the uncertainties. Next, it uses the model from [25] to identify and value private risks. Finally, it uses the procedure for real options analysis from [20] as a valuation model to compare plan-driven and agile projects. The study performs the analysis at the project level.

Real options analysis was chosen as it has been proven a suitable tool to value investment under uncertainties in different contexts, including information technology investments and agile projects [9, 17, 31, 25]. The binomial method was chosen over the Black–Scholes equation as it is easier to adjust parameters over the option’s life and to explain the valuation results because the framework is transparent [20].

The model from [9] was chosen for mapping the life cycle characteristics to real options for three reasons. First, it provides transparency on the source of the flexibility within the life cycles. This makes the comparison between life cycles traceable and easy to understand, and explaining the source of flexibility is straightforward. Second, it differentiates the location of the option relative to the project. The mechanisms occur “in” the project and the types occur “on” the project. Third, it highlights which aspect of the life cycle triggers or contributes to which real option.

The model from Chen et al. [25] was used to evaluate private risks as it provides a relevant risk framework and a valuation model for private risks. These models were selected as they are consistent with the scope of this study. A similar valuation model has been used in evaluating agile projects using real options by [17].

ISO 21502:2020-12, Project, Programme and Portfolio Management—Guidance on Project Management 1st Edition [32] was used as a guideline for defining project roles and processes. It was chosen for two main reasons. First, it was

updated in 2020 to reflect various types of project lifecycles, including agile. Second, it provides a comprehensive view of how projects fit into an organization and guidelines for how a project should be managed. Thus, it was a suitable guide for mapping the project acts and their value drivers.

B. Approach

The aim of the study is to analyze the mechanisms agile projects use to deliver target benefits. To accomplish this goal, the study explores different aspects of project work and maps that understanding to the financial model. Therefore, the proposed approach involves investigating literature, mapping project characteristics to the real options model, and analyzing the results. The proposed procedure was divided into four steps, which are described in the following sections.

- 1) Define the project framework.
- 2) Define the real option characteristics.
- 3) Map the real options characteristics to the project framework.
- 4) Apply the model to an illustrative comparison case.

C. Validity

This approach to defining the model and its elements relies on secondary sources. It is consistent with similar studies [8, 33, 17, 17, 2]. The results were cross-validated by using an illustrative case study.

IV. PROJECT FRAMEWORKS

A. Overview of Project Management Frameworks

The project management bodies of knowledge are used worldwide as the guiding frameworks for standard practices. The guidelines include, for example, *ISO 21502:2020-12, Project, Programme and Portfolio Management—Guidance on Project Management 1st Edition* [32], *APM Body of Knowledge 7th Edition* [11], *A Guide to the Project Management Body of Knowledge (PMBOK guide)—7th Edition* [16], and *Managing Successful Projects with Prince2* [34]. Although criticized by some researchers, the “standards have come to represent an institutionalized collective identity of project managers” [35, p. 37]. In the most recent releases, the guidelines have evolved to suit different delivery approaches: “predictive, incremental, iterative, adaptive or hybrid, including agile approaches” [32, p. 1].

Table II compares the main difference between agile and plan-driven projects.

B. Project Management Processes

The project management standards use various names to describe similar content, processes, and artifacts. Nevertheless, they all cover similar topics on planning, initiating, directing, monitoring, controlling, and closing projects. For the purposes of defining real option mechanisms, we used the subject areas from the ISO standard. The standard provides a holistic and structured way to consider uncertainty and variability in agile and plan-driven projects.

Table III includes an overview of the project roles and processes.

TABLE II
PLAN-DRIVEN AND AGILE COMPARISON

Plan-driven	Agile
One-time delivery of the product as a big release	Frequent releases of the product in increments (usually less than monthly)
Specifications and details exchanged through written documents	Specifications and details exchanged through collaborative communications
Fixed requirements with detailed upfront planning	Initial set of requirements at the start that are subsequently iteratively defined
Detailed plan fixed at the start of the project; formal process to change the plan	High-level plan at the start, adapted through iterative planning
Customer feedback at scheduled sessions (e.g., workshop at the start, follow-up status, validation at end of the project.	Customer and developer feedback almost daily; co-located virtually or physically

TABLE III
PROJECT MANAGEMENT PROCESSES (ISO 21502:2020)

Role	Process	Description
Spr org	Pre-project	Identifying and formalizing the needs and opportunities to realize benefits before authorizing the project.
	Overseeing	Monitoring that the project meets the organizational needs and stakeholder expectations at an acceptable risk level.
	Post-project	Verify the outcomes are sustainable and expected benefits have been realized.
Prj spr	Directing	Engaging directly and through boards to confirm that organizational resources are used as expected or terminating when support is no longer justified.
Pj mgr	Initiating	Plan the project, organization, governance, and management structure and mobilize the project team.
	Controlling	Monitor and measure project performance against agreed plans or authorize changes.
	Closing or terminating	Confirm the completion of the project scope, enable post-project benefits realization, and demobilize resources and facilities.
Work pkg leader	Managing delivery	Define and plan required outputs and outcomes and deliver outcomes to achieve and realize expected benefits.

Legend: Org-Organization, Pkg-Package, Prj-Project, Spr-Sponsor, Mgr-Manager

C. Agile Practices

Each agile framework provides a set of practices and approaches that implement the values and principles defined by the Agile Manifesto. “This results in a high number of agile practices with many variants used in practice and described in literature” [36, p. 1]. For example, the Agile Alliance’s glossary lists 75 different practices. In this section, we identified some of the most frequently used agile practices (in italics) and aligned them in project management subjects from the ISO standards [36, 33].

a) Planning: Breaking requirements into small units called epics or *user stories* that can be *prioritized* and estimated; the prioritized list is called a *backlog*. Deferring decisions on which items to develop until the last moment

and prioritizing requirements at the start of any iteration in an *iteration planning* session [37, 17].

b) Schedule: Delivering product increments in fixed duration iterations (usually two to four weeks). Typically, the objective is to deliver a fully functional product with each release composed of one or more iterations [38].

c) Costs: Estimations in group sessions by the team that will perform the work using wide-band estimation techniques; the estimation process is called *agile estimation* or planning game.

d) Resources: *Self-organizing teams* assign tasks amongst themselves, participate in making decisions, and resolve problems and conflicts. Resources such as people, facilities, equipment, materials, infrastructure, and tools are usually locked in for each iteration.

e) Stakeholders & Communications: Incorporating stakeholders in the project in the iteration planning sessions, *review meetings*, and *daily stand-ups*.

f) Quality: Defining a *definition of done*, performing *retrospectives*, and creating *simple designs* that are part of the quality concepts.

g) Procurement: Contracts where the scope is not predetermined allow features to be implemented based on choices made during the project [17].

h) Lessons learned: *Retrospectives* after each iteration for lessons-learned sessions and improving project performance.

i) Change control: *Small releases* introduce decision points and opportunities to change course using the *iteration planning* process to prioritize the user stories in the backlog [17].

j) Termination: *Small releases* allow the customer to continue or terminate the project at the end of each release [17].

V. REAL OPTION CHARACTERISTICS

A. Real Option Mechanisms

For this study, we define the characterization of real options based on [9]. The mechanisms are the source of flexibility because they are the actions, decisions, or entities that enable a real option. This section describes the mechanisms relevant in a project context when building complex products or processes. The mechanisms described can be applied to agile or plan-driven projects.

1) (Re-)investments: The initial (*re-)investment* mechanism is at work when the sponsoring organization decides to invest in the project by exercising an option. “...the firm has the right to access all cash flows when the project passes the last phase at the investment cost (strike price)” [39, p. 389]. The sponsoring organization can decide to continue or terminate the investment at various stages based on budget, scope, or schedule changes.

2) Iterative Contracting: *Iterative contracting* is organizing contract terms so that resources are allocated one or multiple iterations at a time, and early termination does not result in a

contract breach. The contracting presumes that, where possible, resources are allocated incrementally without negatively impacting quality or creating a moral hazard. A moral hazard is defined as the ineffectiveness or the abuse of trust created by opportunistic behavior [40]. Moral hazards can occur when the supplier and client are from differing organizations. It includes behaviors such as creating a high transparency gap by hiding information, engaging specialists with little competence and experience, or intentionally completing tasks poorly, for example. Moral hazards can be created when the supplier introduces a risk of project failure and increases the transaction cost for the client.

3) *Early Termination*: Project termination describes when a project should end. The factors that cause early termination include significant financial and non-financial losses, high-risk investments, lack of exploitable knowledge created, inability to leverage its exploration and exploitation experience, the collaboration structure, the firm's position in an inter-firm knowledge network, organizational agility, and various project characteristics [41, 42].

An objective criterion for *early termination* is when "the value of benefits (quantifiable or non-quantifiable) does not justify the cost to complete the project, or a more cost-effective alternative is available" [43, p.7].

Early project termination can affect contractual relationships; it implies the contract ends before its expiration date due to force majeure, default by a party, or voluntary buyout [44]. Prerequisites for early termination include termination decision criteria that define the timing and factual basis for the termination [42] and negotiation of contract clauses, dispute resolution processes, or termination compensation.

4) *Modularity*: Modularity is "building a complex product or process from smaller subsystems that can be designed independently yet function together as a whole" [45, p. 1]. *Modularity in design* and *Modularity in use* are two types of real options mechanisms.

a) *Modularity in design*: means dividing the design into visible and hidden components. The visible components follow a set of design rules that consider the architecture, interfaces, and standards [46, 47, 48]. The hidden components may be autonomously designed and developed according to the design rules. Realizing modularity requires architects and designers with a deep level of knowledge to predict future trends and strong leadership to integrate (decentralized) independent teams. Architects compete by attracting designers and having prevalent architecture on the market. Module builders master the production based on hidden information.

In addition, architecture, design rules, interfaces, and standards are key components that enable modularity. The architecture specifies the design's components and controls the visual design rules. The interface describes the interaction, fit, and communication between the subcomponents. The standards establish the tests for conformity to the design rules and measures of the subcomponents' performance. A central architecture requirement is that a hierarchical system can be decomposed into independent, interrelated components and

that stable intermediate forms allow for the rapid evolution of complex systems [49]. One of the design challenges is to decompose the system into subsystems with minimal interdependence between the subsystems [46].

b) *Modularity in use*: allows the end user to compose the final product by mixing and matching components. It accelerates the pace of change, heightens competitive pressures and uncertainty, and transforms companies and markets.

5) *Concurrent Engineering*: Concurrent engineering is the parallel design and development of a product or subsystem. *Point-based* or *setbased concurrent* engineering varies depending on when alternatives converge to the final solution [50]. Waiting longer, developing multiple alternatives, and gathering more information beyond a certain point increases managerial confidence but does not necessarily increase decision-making. Teams waiting longer to eliminate design alternatives gain access to more accurate quality information and make wiser choices on average. Quality does not increase continuously, indicating a transitory period of development work that does not yield additional useful information. Delaying the start of convergence allows the developers to choose any design alternative as new information becomes available [50]. Thus, concurrent engineering can impact development time and product cost and quality.

a) *Point-based concurrent engineering*: various alternatives are available, and the best option is quickly selected to reduce complexity and constrain development costs.

b) *Set-based concurrent engineering*: a range of viable alternatives is selected and eliminated gradually based on weakness to converge at a final solution.

6) *Incremental, Prioritized Delivery*: Incremental, prioritized delivery involves prioritizing feature development based on value and delivering an incremental release for productive use by the sponsoring organization. Erdogmus and Favaro [17] demonstrated in a mathematical study the beneficial effect of the in-project decisions that had the flexibility to prioritize feature development into incremental releases.

B. Real Option Locations

Real option mechanisms in projects can occur "in" the project or "on" projects [9, 30]. Mechanisms "in" projects are the decisions taken by the project team or project manager and occur in the project management processes managed by those roles. These options affect topics such as the system design. Real options "on" projects affect the topic project and its outcome. These options occur in other project management processes, such as those managed by the project sponsor or sponsoring organization.

C. Real Option Types

Real option types are actions or decisions that the owner of the real option can exercise, so they represent flexibility [9]. They can be exercised before the expiration date of an option, after or as uncertainties are clarified. The types are described in Table I.

Fig. 1 demonstrates how an option can be exercised. At the start of the project, point A in the figure, a real option mechanism is implemented. The decision to exercise a real option type associated with that mechanism or not is made at each decision point, points B, C, D, and E in the figure. If the figure represented an agile project, the boxes would be iterations, and if it were a plan-driven project, the boxes would be project phases.

For example, if the real option mechanism at point A was "early termination", the project contract would have a structure that does not penalize the decision for early termination. Then, the real option type "abandon" or "switch" could be exercised at any decision point. For an agile project, a working solution would be expected at the end of each iteration. For a plan-driven project, a working solution would be expected at the last phase.

TABLE IV
REAL OPTIONS MECHANISMS

Mechanism	Description	Location
(Re-) investment	Invest in the project or decide to continue or terminate the investment at various stages based on budget, scope, or time changes.	on
Iterative contracting	Organizing contract terms so that resources are allocated one or multiple iterations at a time, and early termination does not result in a contract breach.	on and in
Early termination	Ending the project when "the value of benefits (quantifiable or non-quantifiable) does not justify the cost to complete the project, or a more cost-effective alternative is available" [43, p.7].	on
Modularity in design	Dividing the design into visible and hidden components	in
Modularity in use	End user to compose the final product by mixing and matching components.	in
Point-based concurrent engineering	A alternative is quickly selected from a range of options to reduce complexity and constrain development costs.	in
Set-based concurrent engineering	A range of viable alternatives is selected and gradually eliminated based on weakness to converge at a final solution.	in
Incremental, prioritized delivery	Prioritizing feature development based on value and delivering an incremental release for productive use by the sponsoring organization.	in

D. Application of Real Options

Table I maps the real options mechanisms, types, and locations. Each row in the table corresponds to a mechanism that can be applied to the project life cycle. This conceptualization is taken from [9], describing real options characterized as a tuple of <Mechanism, Type>.

Fig. 1 provides a visual example of the concept. As an illustration, if the *early termination* mechanism is implemented at time t_0 , then the real options types for *switching* and *abandoning* are enabled. Thus, the sponsor or sponsoring organization could exercise the option to redeploy the project resources for an alternative purpose or abandon the project altogether. Similarly, if the *modularity* mechanism is enabled, it means there is strong team leadership, architecture, and a

design team that allows for modularity in design and use. Thus, the real option to *stage* by partially deploying the system could be exercised to gain early benefits.

E. Private Risks

Projects are affected by uncertainty internal to the project organization and those external produced by market conditions and competitive environments. The internal project-specific risks are private risks [17]. External risks are market and competitive uncertainties independent of the private risks. External risks are usually priced into the financial markets [17]. Private risks can impact the project costs and the market value of product output [22].

The private risks impact project valuations through their unbiased estimates of cash flows. "An unbiased estimate of a cash flow is calculated as a statistical expectation by considering as many scenarios as is feasible and the respective likelihood of these scenarios" [17, p. 8]. Because budget and schedule risks are unique to project situations, they are private.

Historical data, such as risk registers from similar projects, are sources for identifying private risks. A risk register is a structured record, list, or document that details the identified uncertainties and their characteristics [16, 32, 11]. Alternatively, there is a robust body of research on project, information technology, and software development risk factors [51].

Table V is an example of private risks based on a combination of risk factors from [51] and [25].

TABLE V
PRIVATE RISK FACTORS

Risk category	Risk description
Organizational context	Management stability, organizational support for an investment
Project strategy	The adequacy and accuracy of the planning for the project and of the managing and monitoring to ensure that the project reaches its desired goals.
Project structure	The strategic orientation of the application, the number of departments involved, or the business process needs to be changed frequently.
Customers and users	Lack of user involvement during system development, unfavorable attitudes of users towards a new system
Team	Insufficient knowledge or inadequate experience among team members, frequent team member turnover
Technology complexity	Whether the new technology is used, the complexity of the processes being automated, whether many links to existing systems are required
Changing requirements	Frequently changing requirements; incorrect, unclear, inadequate, or ambiguous requirements

F. Applying Real Options Binomial Tree

The next step is to apply the discussed components to a project analysis. The steps include defining the project use case by defining the real option tuple, quantifying private risks, identifying the input parameters, calculating the option parameters, building the binomial tree, calculating the asset

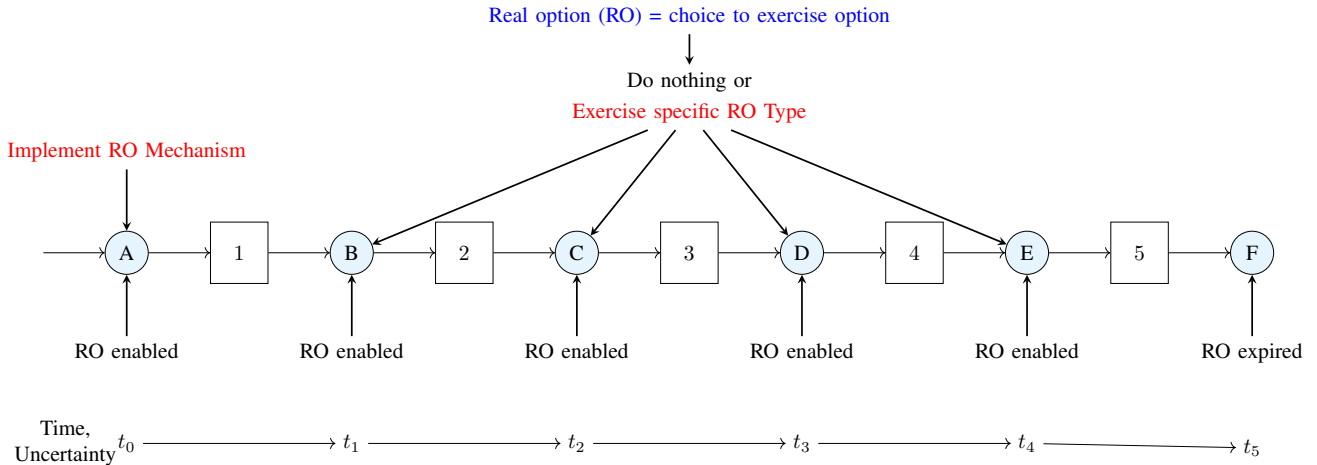


Fig. 1. Real Options Example

values at each tree node, and analyzing the results. The following sections describe the inputs and computations, closely following those given by [20].

1) *Define the use case:* For a given project use, the mechanism applicable to the project should be determined and then mapped to the real option types to form a tuple of <Mechanism, Type>. The real option types determine the quantitative parameter to compute the real option put values.

2) *Quantify the private risks:* The private risk factor (F) can be applied as a risk premium or failure rate to the uncertainty model. The risk premium is a quantitative assessment of the financial impacts to address the risk. The value of the risk premium is added to the project investment cost. It measures additional exposure to the risk, and the sensitivity measure is the degree of exposure. “The failure rate describes the extent to which the IS [Information Systems] project will be a failure” [25, p. 779]

Here, we follow the method of using failure rate as described and proposed by [25] because it considers overall project risks. The scores of the probability (P_f) and consequences (C_f) of the risks are from low to major as follows: 0.1-low, 0.2-minor, 0.5-moderate, 0.7-significant, and 0.9-major.

“The scores for each individual dimension of probability (P_f) and consequence (C_f) are added and the sum is divided by the number of factors used to assess them:

$$P_f = \sum P_i/i \quad C_f = \sum C_i/i \quad (1)$$

The formula for the failure rate is as follows in Equation 2:

$$F = P_f + C_f - (P_f)(C_f) \quad (2)$$

” [25, p. 779].

3) *Identify input valuation parameters:*

a) *Underlying asset value (S_o):* The underlying asset value (S_o) is based on the discounted cash flow of the expected free cash flows, and “production phase free cash flows are the

net revenues calculated from the expected future revenues and costs associated with project output in its production phase” [20, p. 72]. It is an estimate based on an assumption of the value of the project outcome.

b) *Exercise price (X):* The exercise price (X) is the investment cost. Investment costs are the development cost and production phase capital costs. They are primarily influenced by how accurately the organization estimates the costs and how efficiently the project is brought to completion; these costs are influenced by private risks and are not influenced by market forces [20]. The investment costs directly impact the option value.

$$X = X * R_p \quad (3)$$

c) *Option life (T):* The option life parameter (T) is the time to maturity. Unlike financial options that have a known maturity date, in real options, the option life is estimated based on the expected duration of the project’s development and production phases. Factors such as loss of market share and competition can reduce option value even as maturity time increases. Thus, the option life should be established sufficiently long to clear uncertainty but not so long as to reduce the value of the project outcome.

d) *Volatility of underlying asset (σ):* The volatility factor (σ) is the measure of the variability of the underlying asset over its lifetime. It is given by the standard deviation of the continuous rate of return on the asset value over time. The volatility factor varies in step with the time step.

In a project context, using historical data from similar projects with similar market performance and cash flow profiles is an acceptable method of accounting for market reality [20].

e) *Time step (δt):* The time step factor (δt) splits the option’s life into time increments. “The higher the time increments used in the binomial method, the closer you will get to this value [Black–Scholes equation] ...with only four to six

TABLE VI
REAL OPTIONS MECHANISMS AND TYPE

Real Option Mechanism "on" the project	Real Option Mechanism "in" the project	Types	Real Option Meaning	Exercise Price (X value) Parameter
Early termination	Incremental, prioritized delivery	Switch	Within the project, you have the option to prioritize deliveries.	Investment costs
(Re-)investment		Change scale: expand	Continue to operation and keep the option open for the future or exercise one of the options.	Expansion factor, cost of expansion
Iterative contracting	Iterative contracting	Change scale: reduce	To maximize continue as planned or reduce planned expenses by the contract factor and save the savings values. At each node, you have an option to continue the operation and keep the option open or to contract by half.	Contract factor, savings of contracting
	Modularity, Incremental, prioritized delivery	Stage	Divide project into multiple investment packages.	Investment costs
Early termination	Modularity	Abandon	At each node, you have the option to either abandon the project for a salvage value or continue until the option expires. Each node represents the value maximization of abandonment versus continuation.	Salvage value
(Re-)investment	Set-based concurrent engineering	Strategic growth	Each node represents the value maximization of continuation versus expansion. At each node, you can either continue the operation and keep the option open for the future or expand it by committing the investment for expansion.	Expansion factor, cost of expansion
	Modularity (in use) Incremental, prioritized delivery, Point-based concurrent engineering	Defer with revenue leakage	The annual leakage rate or revenue lost due to the delayed investment	Investment or annual leakage or revenues lost due to the delayed investment.
	Modularity (in use), Incremental, prioritized delivery	Defer	Option to invest in production development or the option to wait until the next time period before the option expires.	Investment costs

time steps, a relatively good approximation can be obtained" [20, p. 145]. "The higher the number of time steps, the higher the level of granularity and therefore the higher the level of accuracy of option valuation" [20, p. 113].

f) *Risk-free interest rate (r)*: The risk-free interest rate parameter (r) is the current interest rate on the risk-free asset. In a "real options model [it] is usually determined based on the U.S. Treasury spot rate of return, with the maturity equivalent to the option's time to maturity" [20].

4) *Calculate the option parameters*: The uncertainty model is defined using the binomial tree. The tree is built by starting with the underlying asset value S_o as the first node on the left and multiplying it by the up factor (u) for the up node and the down factor (d) for the down node for the first step. Then, moving to the right, perform this action for each node in the binomial tree for each time step until the last time step. The results are the expected asset values at each node of the tree. The up and down factors are based on the volatility of the underlying asset S_o and calculated using Equations 4 and 5.

$$u = \exp(\sigma\sqrt{\delta t}) \quad (4)$$

$$d = 1/u \quad (5)$$

The option values are calculated using the risk-neutral probabilities (p) as represented by Equation 6. It represents the riskless asset during the life of the option. The model involves "risk adjusting the cash flows throughout the lattice with risk-

neutral probabilities and discounting them at the risk-free rate" [20, p. 115].

$$p = \frac{\exp(r\delta t) - d}{u - d} \quad (6)$$

5) *Defining the uncertainty model*: Build the binomial tree by using the time step (δt) for the options life (T) to calculate the asset values over the life of the option, starting at the underlying asset value (S_o) and multiplying by the up factor (u) and the down factor (d) at each time step, moving right until the last time step. The value computed is the asset value at each node.

The option values calculated at each node will depend on the real option type chosen. For example, if the option type is to abandon the project, then at each node, the option price will be based on comparing the exercise price (X) to the value of keeping the option open until it expires. Starting with the last node, compare the asset value on the node as computed above with the exercise price (X), and to maximize the value, the biggest value is the option value. As an illustration, if the asset value at the node is 100 USD and the exercise price is 10 USD, then the option value would be 100 USD.

At the intermediate nodes, those steps away from the last node, "the expected asset value for keeping the option open...is the discounted (at the risk-free rate) weighted average of potential future option values using the risk-neutral probability as weights:"

$$[p(S_o u^5) + (1 - p)(S_o u^4 d)] * \exp(-r\delta t) \quad (7)$$

[20, p. 159]. The formula is an example of computing the asset value at node S_0u^4 .

The option valuation binomial tree is computed backward to time zero, and the higher of the exercise price or computed expected asset value is retained.

6) *Analyzing the results:* After quantifying relative benefits and costs as relative option values at each tree node by backward induction, the retained put option value at time zero is the real option put value. Kodukula and Papudesu [20] notes this real option put value should be a good approximation to that produced using the Black–Scholes equations.

a) *Probability of success using the failure rate:* The value at each node would consider the impact of the risk failure rate by educing the chances of achieving the value [25].

$$[\max[0, (1 - F)S_0u^4 - I_4]] \quad (8)$$

The formula is an example of calculating the value at node S_0u^4 using the project investment (I) in the time period (t) four and the failure rate (F).

b) *Probability of success from the binomial lattice:* The information in the binomial tree has further uses, such as using the probabilities for investment decisions. For example, summing the number of steps in the binomial lattice, computing paths where the options will not be exercised, and dividing the two gives a probability that the options will not be profitable. This value can be used to compute the probability that the project will succeed. For example, there are 32 paths for all end nodes in a five-step lattice. If the two bottom-end nodes where the option will not be exercised include six paths, then the probability of failure will be 6/32 or 19%. Thus, the probability of success will be 81%. See [20] for more information.

VI. ILLUSTRATIVE CASE COMPARISON

A. Case Background

The project is an illustrative case using the abandon real option type to compare the agile and plan-driven methodologies. The project entails bringing a new product to the market. The total estimated cost to bring the product to market, including development, is \$95 million, and it should return \$100 million at a market risk-adjusted discount rate. The project is estimated to last five years. The \$95 million investment cost is distributed evenly at \$19 million per year. If the project can deliver a minimally viable product to the market, it would have a payout of \$65 million; we refer to this as salvage value. The annual volatility rate of the future cash flows is 35%, and the annual risk-free interest rate of the period is 5%. The data in this example follow the structure of the simple example from [20].

B. Define the use case

Referring to Table VI, for the abandon option, the tuple of <Mechanism, Type> suggests that the sponsoring organization or project sponsor should have organized the contractual model for early termination. Further, the project team could influence the ability to deliver value prior to abandonment if they have included modularity in the design and use.

TABLE VII
ILLUSTRATIVE CASE RISK SCORES

ID	Risk Category	Plan-driven		Agile	
		P_f	C_f	P_f	C_f
R1	Organizational context	0.28	0.52	0.28	0.52
R2	Project strategy	0.00	0.00	0.00	0.00
R3	Project structure	0.2	0.30	0.2	0.20
R4	Customers and users	0.18	0.22	0.1	0.20
R5	Team	0.10	0.18	0.1	0.50
R6	Technology complexity	0.42	0.60	0.42	0.60
R7	Changing requirement	0.3	0.32	0.3	0.10

TABLE VIII
ILLUSTRATIVE CASE FAILURE RATES

Risk failure rate	F (agile) = 0.50
Risk failure rate	F (plan-driven) = 0.52

C. Quantify the Private Risks

The private risk factor (F) varies per methodology type and is based on the risk options mechanisms. For this study, we used the risk scores from Chen et al. [25] for the plan-driven project and adjusted the parameter for the agile project.

First, the modularity mechanisms discuss the need for strong leadership, architecture, and design skills. Thus, the lack of these capabilities would severely impact an agile project. Next, the delivery method for agile projects suggests high customer engagement, so the probability of this risk would be low. The probability of many requirement changes would be similar for both project types. However, the consequences in an agile project would be lower than in a plan-driven one. Therefore, the private risk profile for the two types of projects would differ.

The corresponding risk scores are shown in Table VII, and the risk failure rates are shown in Table VIII.

D. Applying Real Options Binomial Tree

Based on the inputs discussed, the next step is to use the binomial tree to compute the real options based on the uncertainty model. The following sections describe the inputs and computations.

1) *Identify valuation parameters:* The input parameters for the real options calculation are provided in Table IX and are the same for agile and plan-driven projects, except for the Exercise price, which represents the abandoned- value for the delivery of a minimally viable product after early termination of the project.

TABLE IX
ILLUSTRATIVE CASE INPUT PARAMETERS

Project investment	I = \$95 million
Underlying asset value	S_0 = \$100 million
Exercise Price	X = \$65 million
Option Life	T = 5 years
Volatility of underlying asset	σ = 35%
Time Step	δt = 1 year
Risk-free interest rate	r = 5%

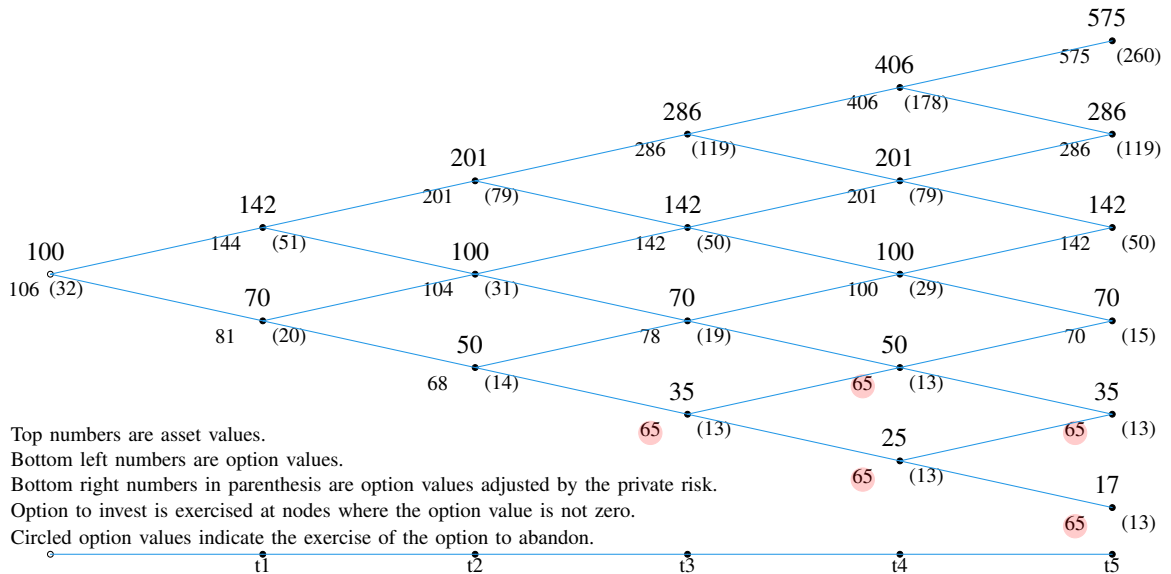


Fig. 2. Binomial Tree for Agile Project with an Option to Abandon




Fig. 3. Binomial Tree for Option for Plan-Driven Project with no Abandon Option

TABLE X
ILLUSTRATIVE CASE OPTION PARAMETERS

Up factor	$u = 1.42$
Down factor	$d = 0.70$
Risk-neutral probabilities	$p = 0.49$

2) *Calculate the option parameters:* The same formula calculates the option parameters for both methodology types, and the values are shown in Table X.

3) *Defining the uncertainty model:* The binomial tree is shown in Fig. 2 is for the agile project and Fig. 3 is for the plan-driven project. They show a binomial lattice where the top numbers are asset values, and the bottom left are option values. The bottom right number is the risk-adjusted option values. From the left, it starts at time = 0 with the underlying asset value (S_0) and real option project value just below. The paths proceed to the right with the asset and option values for the five years.

4) *Analyzing the results:* The binomial tree quantifies the relative benefits and costs at each tree node. First, if the project succeeds as planned for the full five years, the agile project would yield \$106 million, considering the option to recover if abandoned. The plan-driven project does not have the abandon option; thus, it would yield \$100 million. Both the agile and plan-driven projects deliver a respectable return. For agile, it would be \$11 million based on a \$95 million investment and \$5 million for the plan-driven.

Second, for our purpose of comparing agile and plan-driven methods, if the project ends earlier than five years, then the ability to benefit at all would depend on whether the project delivered a minimally viable product in the year before abandonment. Because we used the modularity mechanisms, at a minimum, an agile project would recover the salvage value of \$65 million before we consider the failure rate. If we pursued the plan-driven project and stopped before the final delivery, then the salvage value would also be at risk. Thus, the option values are lower starting at year 3.

Finally, considering the failure rate, the risk failure rate must reduce the option value and subtract the investment costs. For example, at the beginning of year 5, the $S_0 u^5$ option value of \$575 million would be \$285 million for an agile project and \$260 million for a plan-driven project using Equation 8. Further, the S_0 real option value would be \$32 million for the agile project and \$29 million for the plan-driven project.

VII. CONCLUSIONS

As noted by Kodukula and Papudesu [20, p. 95], “ROA [Real Options Analysis] is most valuable when there is high uncertainty with the underlying asset value and management has significant flexibility to change the course of the project in a favorable direction and is willing to exercise the options.” In this study, we identified the mechanisms that provide management flexibility to act and change the course of agile or plan-driven projects.

First, the sponsoring organization and project sponsor can use investment and termination options to deploy corporate

resources away from a project. Next, the project manager and the project team can make in-project decisions that enable or prohibit certain investment options. For example, a project that does not have a modular design or prioritizes iterative and incremental delivery inhibits deriving value from an early termination.

Second, the study relates project acts (e.g., mechanism) to the real options types and computation model. The binomial tree provides a visual method for connecting a project timeline to the economic value of the expected benefits. Further, the project failure used in the computation explains the consequences of the risks on the economic value. Thus, the research provides a multi-faceted view for evaluating a project’s contribution to organizational benefits. Finally, the analysis structure compares how different project approaches react under similar uncertainty scenarios.

A. Contributions to Knowledge

The results quantify subjective and theoretical speculation on how agile projects contribute to realizing organizational benefits. Using real options, the study connects the project structure to the methods for calculating organizational benefits. The results contribute to the project management literature on agile methodologies and the theory of benefits realization in projects.

B. Implications for Practice

Three management levels are involved in the governance and structuring of project work. The study offers some guidance for people at those levels to understand the mechanisms that make an agile project function and how those mechanisms can be used to contribute to organizational benefits. The sponsoring organizations, sponsors, and project managers should consider ways to structure project teams, contracts, and investments to retain the flexibility to deploy or redeploy corporate resources. The project managers and work package leaders should consider how the internal operations of the project should be organized to allow for changes in corporate direction.

C. Implications for Research

In future studies, researchers investigating agile methodologies should consider methods for investigating and validating the mechanisms that contribute to project performance. Most research simply accepts that agile projects improve project performance. However, there is insufficient empirical evidence on the mechanisms that explain performance. Further, the risk failure rate for agile and plan-driven methodologies differs slightly, for example, by a few points in our illustrative case. However, the time dimension of the project is the strongest indicator of achieving organizational benefits. The longer the project, the more likely it is to fail. This aspect was not considered sufficiently in the modeling of this work, and it offers an interesting and important area for future research.

D. Limitations and Further Research

The results of this study are not generalizable beyond information technology projects due to the type of mechanisms identified in this research. Specifically, software development projects have been the most active in applying agile methodologies. No steps were designed to determine whether the proposed methods would apply to other project types. Further, we lacked financial or factual data to comprehensively measure an actual case study.

REFERENCES

- [1] O. Zwikael, Y.-Y. Chih, and J. R. Meredith, "Project benefit management: Setting effective target benefits," *International Journal of Project Management*, vol. 36, no. 4, pp. 650–658, 2018.
- [2] C. Marnewick and A. L. Marnewick, "Benefits realisation in an agile environment," *International Journal of Project Management*, vol. 40, no. 4, pp. 454–465, 2022.
- [3] P. Serrador and J. K. Pinto, "Does agile work? — a quantitative analysis of agile project success," *International Journal of Project Management*, vol. 33, no. 5, pp. 1040–1051, 2015.
- [4] A. Gemino, B. Horner Reich, and P. M. Serrador, "Agile, traditional, and hybrid approaches to project success: Is hybrid a poor second choice?" *Project Management Journal*, vol. 52, no. 2, pp. 161–175, 2021.
- [5] R. A. Lundin and A. Söderholm, "A theory of temporary organization," *Scandinavian Journal of Management*, vol. 11, no. 4, pp. 437–455, 1995.
- [6] K. Beck, M. Beedle, A. v. Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, J. Kern, B. Marick, R. C. Martin, S. Mellor, K. Schwaber, J. Sutherland, and D. Thomas, "Manifesto for agile software development," Dec 2001.
- [7] J. Binder, L. I. V. Aillaud, and L. Schilli, "The project management cocktail model: An approach for balancing agile and ISO 21500," *Procedia Social and Behavioral Sciences*, vol. 119, pp. 182–191, 2014.
- [8] D. F. Rico, "What is the roi of agile vs. traditional methods? an analysis of xp, tdd, pair programming, and scrum (using real options)," *unpublished*, 2008.
- [9] T. Mikaelian, D. J. Nightingale, D. H. Rhodes, and D. E. Hastings, "Real options in enterprise architecture: A holistic mapping of mechanisms and types for uncertainty management," *IEEE Transactions on Engineering Management*, vol. 58, no. 3, pp. 457–470, 2011.
- [10] R. Gunther McGrath and A. Nerkar, "Real options reasoning and a new look at the R&D investment strategies of pharmaceutical firms," *Strategic Management Journal*, vol. 25, no. 1, pp. 1–21, 2004.
- [11] APM, "APM body of knowledge seventh edition," 2019.
- [12] T. Cooke-Davies, "The "real" success factors on projects," *International Journal of Project Management*, vol. 20, no. 3, pp. 185–190, 2002.
- [13] R. J. Turner, "Projects for shareholder value the influence of projects at different financial ratios," in *29th Annual Project Management Institute 1998 Seminars & Symposium*. Project Management Institute, 1998, Conference Proceedings.
- [14] P. Serrador and R. Turner, "The relationship between project success and project efficiency," *Project Management Journal*, vol. 46, no. 1, pp. 30–39, 2015.
- [15] R. Joslin and R. Müller, "The impact of project methodologies on project success in different project environments," *International Journal of Managing Projects in Business*, vol. 9, no. 2, pp. 364–388, 2016.
- [16] PMI, *A Guide to the Project Management Body of Knowledge (PMBOK® Guide)—Seventh Edition*. Newtown Square, PA: Project Management Institute, 2021.
- [17] H. Erdogmus and J. Favaro, "Keep your options open: Extreme programming and the economics of flexibility," in *Giancarlo Succi, James Donovan Wells and Laurie Williams, "Extreme Programming Perspectives"*, Addison Wesley, 2002, 2002, Conference Paper.
- [18] Y. Shastri, R. Hoda, and R. Amor, "Does the 'project manager' still exist in agile software development projects?" in *2016 23rd Asia-Pacific Software Engineering Conference (APSEC)*, Hamilton, New Zealand, 2016, Conference Paper, pp. 57–64.
- [19] A. Tiwana, M. Keil, and R. G. Fichman, "Information systems project continuation in escalation situations: A real options model," *Decision Sciences*, vol. 37, no. 3, pp. 357–391, 2006.
- [20] P. Kodukula and C. Papudesu, *Project valuation using real options: A practitioner's guide*. J. Ross Publishing, 2006.
- [21] R. G. Fichman, M. Keil, and A. Tiwana, "Beyond valuation: "options thinking" in IT project management," *California Management Review*, vol. 47, no. 2, pp. 74–96, 2005.
- [22] M. Benaroch and J. Goldstein, "An integrative economic optimization approach to systems development risk management," *IEEE Transactions on Software Engineering*, vol. 35, no. 5, pp. 638–653, 2009.
- [23] F. D. Maddaloni, G. Favato, and R. Vecchiato, "Whether and when to invest in transportation projects: Combining scenarios and real options to manage the uncertainty of costs and benefits," *IEEE Transactions on Engineering Management*, pp. 1–15, 2022.
- [24] J. L. Wellman, "Project valuation using real options: A practitioner's guide," *Project Management Journal*, vol. 37, no. 5, p. 116, 2006.
- [25] T. Chen, J. Zhang, and K.-K. Lai, "An integrated real options evaluating model for information technology projects under multiple risks," *International Journal of Project Management*, vol. 27, no. 8, pp. 776–786, 2009.
- [26] I. Krystallis, G. Locatelli, and N. Murtagh, "Talking about futureproofing: Real options reasoning in complex infrastructure projects," *IEEE Transactions on Engineering Management*, pp. 1–14, 2020.
- [27] C.-H. Wang and K. J. Min, "Electric power generation planning for interrelated projects: a real options

- approach,” *IEEE Transactions on Engineering Management*, vol. 53, no. 2, pp. 312–322, 2006.
- [28] A. Jahanshahi Asghar and A. Brem, “Does real options reasoning support or oppose project performance? empirical evidence from electronic commerce projects,” *Project Management Journal*, vol. 48, no. 4, pp. 39–54, 2017.
- [29] Z. Racheva, M. Daneva, and L. Buglione, “Complementing measurements and real options concepts to support inter-iteration decision-making in agile projects,” in *2008 34th Euromicro Conference Software Engineering and Advanced Applications*, 2008, Conference Proceedings, pp. 457–464.
- [30] T. Wang and R. de Neufville, “Identification of real options ‘in’ projects,” in *Systems Sciences*, vol. 16, 2005, Conference Proceedings, pp. 1124–1133.
- [31] H. Erdogmus, “Valuation of learning options in software development under private and market risk,” *The Engineering Economist*, vol. 47, 2002.
- [32] ISO, “ISO 21502: 2020-12 project, programme and portfolio management — guidance on project management,” 2020.
- [33] G. J. Miller, *Framework for Project Management in Agile Projects: A Quantitative Study*, ser. Information Technology for Management: Current Research and Future Directions. Springer International Publishing, 2020.
- [34] The Stationery Office, “Managing successful projects with PRINCE2,” 2017.
- [35] P. Eskerod and M. Huemann, “Sustainable development and project stakeholder management: What standards say,” *International Journal of Managing Projects in Business*, vol. 6, no. 1, p. 36–50, 2013.
- [36] M. Neumann, “The integrated list of agile practices — a tertiary study,” in *Lean and Agile Software Development*, A. Przybyłek, A. Jarzębowicz, I. Luković, and Y. Y. Ng, Eds. Springer International Publishing, 2022, Conference Proceedings, pp. 19–37.
- [37] A. Appari and M. Benaroch, “Monetary pricing of software development risks: A method and empirical illustration,” *Journal of Systems and Software*, vol. 83, no. 11, pp. 2098–2107, 2010.
- [38] H. Alahyari, R. Bertsson Svensson, and T. Gorschek, “A study of value in agile software development organizations,” *Journal of Systems and Software*, vol. 125, pp. 271–288, 2017.
- [39] M. Montajabiha, K. Alireza Arshadi, and B. Afshar-Nadjafi, “A robust algorithm for project portfolio selection problem using real options valuation,” *International Journal of Managing Projects in Business*, vol. 10, no. 2, pp. 386–403, 2017.
- [40] B. Wachnik, *Moral Hazard in IT Project Completion. An Analysis of Supplier and Client Behavior in Polish and German Enterprises*, ser. Information Technology for Management. Cham: Springer, 2016, vol. 243, pp. 77–90.
- [41] H. Delerue and H. Sicotte, “Resource interdependence and project termination: An analysis in the biopharmaceutical industry,” *International Journal of Project Management*, vol. 38, no. 5, pp. 256–266, 2020.
- [42] M. Vaculík, A. Lorenz, N. Roijakkers, and W. Vanhaverbeke, “Pulling the plug? investigating firm-level drivers of innovation project termination,” *IEEE Transactions on Engineering Management*, vol. 66, no. 2, pp. 180–192, 2019.
- [43] J. Schmidt, “IT project failure, termination and the marginal cost trap,” *Journal of Modern Project Management*, vol. 10, pp. 255–275, 2022.
- [44] W. Xiong and Y. Han, “Incentives of early termination compensation in public–private partnership projects,” *IEEE Transactions on Engineering Management*, pp. 1–13, 2021.
- [45] C. Y. Baldwin and K. B. Clark, “Managing in an age of modularity,” *Harvard Business Review*, vol. 75, no. 5, pp. 84–93, Sep/Oct 1997, copyright - Copyright Harvard Business Review Sep/Oct 1997 Last updated - 2015-11-06.
- [46] E. J. de Waard and E.-H. Kramer, “Tailored task forces: Temporary organizations and modularity,” *International Journal of Project Management*, vol. 26, no. 5, pp. 537–546, 2008.
- [47] R. N. Langlois, “Modularity in technology and organization,” *Journal of Economic Behavior & Organization*, vol. 49, no. 1, pp. 19–37, 2002.
- [48] G. Lizarralde, M. d. Blois, and I. Latunova, “Structuring of temporary multi-organizations: Contingency theory in the building sector,” *Project Management Journal*, vol. 42, no. 4, pp. 19–36, 2011.
- [49] H. A. Simon, “The architecture of complexity,” *Proceedings of the American philosophical society*, vol. 106, no. 6, pp. 467–482, 1962.
- [50] J. K. Liker, D. K. Sobek, A. C. Ward, and J. J. Cristiano, “Involving suppliers in product development in the United States and Japan: evidence for set-based concurrent engineering,” *IEEE Transactions on Engineering Management*, vol. 43, no. 2, pp. 165–178, 1996.
- [51] E. Ziemba and I. Kolasa, “Risk factors relationships for information systems projects – insight from polish public organizations,” in *Information Technology for Management. Lecture Notes in Business Information Processing*, E. Ziemba, Ed., vol. 243. Cham: Springer, 2016, Book Section, pp. 55–76.

MBSPI—A Model-Based Security Pattern Integration Approach for software architectures

Anas Motii
0009-0005-4936-0028
Mohammad VI Polytechnic University
College of Computing
Benguerir, Morocco
Email: anas.motii@um6p.ma

Mahmoud El Hamlaoui
0000-0003-3315-7373
Mohammed V University in Rabat
Rabat, Morocco
Email: mahmoud.elhamlaoui@ensias.um5.ac.ma

Abstract—Incorporating security patterns into software architecture is essential for robust system design. Model Driven Engineering (MDE) offers a structured approach to software development, emphasizing modeling and automation. This paper explores the integration of security patterns into software architecture using MDE techniques. We highlight the benefits of this approach, including improved level of security and enhanced maintainability. Challenges such as modeling complexity and tool support are also discussed. Through a SCADA (Supervisory Control and Data Acquisition) system case study, we demonstrate the effectiveness of integrating security patterns into software architecture using MDE.

Index Terms—Security patterns, Pattern Integration, Software architecture, MDE, OCL

I. INTRODUCTION

IN THE realm of system and software architecture, while there exists a fundamental understanding of security engineering principles, there is often a notable gap in best practices necessary for implementing security measures derived from risk assessments. This gap has propelled the exploration of security patterns as a prominent field of study in recent years. Security patterns serve to encapsulate and disseminate expert knowledge by offering generic, reusable solutions to frequently encountered security challenges, particularly those related to architecture. This methodological approach seeks to equip architects with the tools required to effectively address and mitigate common security vulnerabilities. A complete catalog of security patterns has been introduced by Fernandez [1]. Unfortunately, there are two major issues. First, traditional security patterns are usually described as informal guidelines to solve a certain problem using templates such as POSA (Pattern-Oriented Software Architecture) and GoF (Gangs of Four). Despite the benefits of security patterns in promoting reuse, their practical application faces significant challenges. The disconnection between the threat models identified through risk assessment and the protective measures outlined in security patterns creates a substantial barrier. Additionally, the manual integration of these patterns into architectural designs introduces further complexity, often leading to incorrect implementations. This misalignment not only hampers the effective utilization of security patterns but also leaves critical security issues unresolved, underscoring the need for

improved methodologies in pattern integration and application. One notable finding in [2] is the lack of emphasis on pattern integration, which served as a key motivator for our research endeavors.

Drawing on the advantages of Model-Driven Engineering (MDE) such as improved quality and productivity, we use dedicated modeling languages and Model-To-Model techniques tailored to conduct a comprehensive pattern integration process. In this paper, we introduce a Model-Based Security Pattern Integration (MBSPI) approach for software architecture and its tool support. We use the Object Constraint Language (OCL) for the formalization of security properties.

The remainder of the paper is organized as follows. Section II identifies related work to pattern integration. Section III presents the main steps of the MBSPI approach. The MDE framework is described in section IV, more specifically, Model-to-Model transformations, and OCL constraints. Section V specifies requirements for tool support. In section VI, MBSPI is assessed over a SCADA (Supervisory Control and Data Acquisition) system case study. Finally, section VII, concludes and sums up the contributions and future work.

II. RELATED WORK

Over the years there has been a noticeable divorce between pattern experts and pattern users [3]. On one hand, pattern experts create and document patterns and on the other hand, pattern users are rarely aware of relevant patterns. In addition, the latter do not have a good understanding of how to leverage and apply a pattern. Works relevant to the integration of patterns and aspects are discussed since this issue (i.e., integration) has been tackled in both research areas.

In [4], the authors explained how pattern integration can be achieved by using a library of precisely described and formally verified solutions. In [5], the authors present an approach for creating a security-enhanced system model using the SecFutur Engineering Process and the SecFutur Process Tool (SPT). In [6], the authors introduced a method and tool support for developing secure and private IT systems using Computer Supported Security Patterns (COSSP). The integration process targets object-oriented applications.

Aspect-Oriented Modeling is similar to pattern modeling with regards to encapsulating concerns such as security for use. However, the difference is that aspects are part of software fulfilling a function dealing with the design stage, whereas patterns can deal with different development stages. In [7], Nguyen et al. presented a pattern-driven secure system development process combined with an aspect-oriented security design methodology. To use this approach, the designer is required to manually construct security solutions of the considered system and the definition of mappings of these solutions into the model under development. Mouheb et al. [8] developed a UML profile that allows modeling security mechanisms as UML annotated aspect models to be woven automatically into a UML design model. Horcas et al. [9] propose an Aspect-Oriented Modeling (AOM) approach to weaving customized security models into an application using the Common Variability Language (CVL) and the Atlas Transformation Language (ATL). These works have left the Verification & Validation activity for future work. In addition, conflicts between the design and other architectural attributes may occur during this task after the weaving. In [10], Georg et al. proposed an approach for modeling security mechanisms and attacks as aspects using UML. To prove that the integration is correct, they use model verification on the application composed of the attack model and the security mechanisms.

The authors in [2] have reviewed security pattern specifications and usage. The conclusion indicates that there is minimal attention given to pattern integration. Peldszus et al. [11] introduce the GRaViTY approach which offers tools to align various artifacts generated during the model-driven development process, along with mechanisms for defining, applying, and reusing security requirements. Consequently, while the GRaViTY approach shows promise in supporting model-driven development, its limited scope and lack of empirical evidence may restrict its applicability and effectiveness in real-world scenarios. The work of [12] addresses the critical challenge of ensuring security in online service-oriented systems through a pattern-oriented approach. While numerous security design patterns exist, their integration remains a less explored area, motivating the proposed methodology. By utilizing the algebraic specification language SOFIA (Service-Oriented Formalism In Algebras) and translating specifications into the Alloy formalism, the work introduces a systematic approach to verify the validity and correctness of security design pattern compositions. The development of a tool support facilitates automated verification, demonstrated through a crowdfunding application case study. Despite advancements, existing works lack focus on proving functional correctness of pattern compositions, presenting an opportunity for future research. Additionally, the work highlights the potential of algebraic specifications for automated testing, suggesting avenues for further experimentation and extension of the proposed methodology. The approach faces scalability challenges with complex systems and struggles with adaptability to new security threats. Practicality concerns in environments that favor rapid development may also arise, questioning the method's

efficiency in fast-paced settings. These factors highlight the need for a balance between rigorous security verification processes and the dynamic requirements of modern software development cycles. In [13], the authors introduce a groundbreaking aspect-oriented models-centered security framework aimed at addressing security challenges intrinsic to intelligent systems. The framework's components are carefully crafted based on identified threats specific to intelligent systems, with each element depicted through Unified Modeling Language (UML) diagrams. Other works like Armoush and al. [14] have focused on the integration of security patterns in the design of safety-critical embedded systems.

III. MBSPI APPROACH

In this section, the MBSPI approach is presented. This process is based on the approach done in [15] and previous work [16]–[18] which provides the first solution for design pattern integration in the context of object-oriented applications. Here, we go a step further. The validation of security properties and constraints has been formalized with OCL¹, which is a standard developed by the Object Management Group (OMG).

Fig. 1 depicts the MBSPI process which consists of five phases: Preparation, Elicitation, Context Validation, Merge, and Verification & Validation. The phases are described after presenting the pattern integration artifacts consumed and produced by each phase.

A. Definitions

The process interacts with the following artifacts:

- **Security Pattern** represents a modular part of a system that encapsulates a solution of a recurrent security problem in a specific context. The pattern and its constituents are developed by an security expert.
- **Application Diagram** is the representation of the software architecture of the application.
- **Pattern representation artifact** is the security pattern solution. It represents the architectural solution of the pattern. It consists of a number of software components: *participants* and *security mechanisms*.
- **Participants (roles)** represent generic components of a *security pattern solution*. They are the roles potentially played by components of the application.
- **Security mechanisms** are software components part of the *security pattern solution*. They provide primitive security functions (e.g., encryption, signing). A library of these functions is provided in order to allow security pattern solution modeling.
- **Preconditions** are the constraints that the application must verify in order for the integration to work.
- **Postconditions** are pattern security properties that the application verifies after the integration of the pattern.
- **Casting Diagram** consists of the *application diagram*, the *pattern representation artifact* diagram and the *bindings* between components of the two diagrams. The

¹<https://www.omg.org/spec/OCL/>

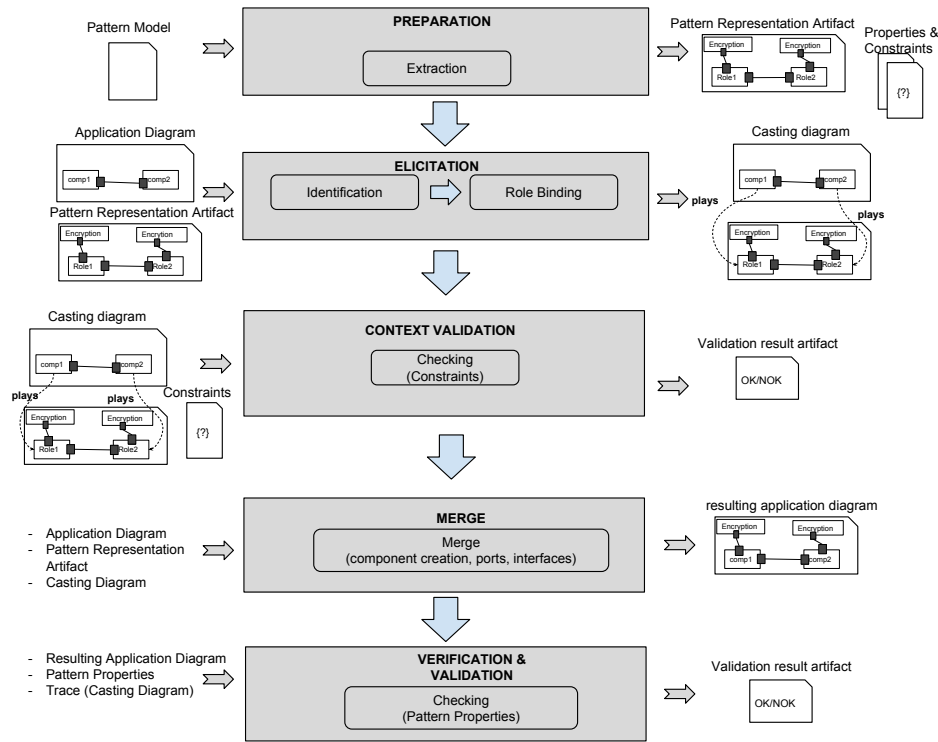


Fig. 1. MBSPI Process

bindings constitute mappings to identify the application components that play roles (participants) in the security pattern solution.

- **Resulting Application Diagram** is the final software architecture diagram once the pattern has been integrated into the application.

B. Hypothesis

Below are a set of assumptions that were used to construct the integration process:

- **Hypothesis 1: Security feature incorporation.** The pattern integration process is only intended to incorporate security features. We verify that the application diagram contains the necessary functional features before starting the integration.

For example, in order to integrate a Secure Communication Pattern between a Client and Server components:

- There must be communication.
- Client and Server interfaces must have *Send()* and *Receive()* interfaces.
- **Hypothesis 2: Pattern properties relevant to messages.** The focus is particularly on the following pattern properties intended for transmitted messages:
 - **Property 1.** Confidentiality of messages.
 - **Property 2.** Integrity of messages.
 - **Property 3.** Authenticity of messages.
- **Hypothesis 3: Pattern constraints.** Preconditions are conditions that the artifacts must meet before going further:

- **Precondition 1.** All pattern participants should be bound (one participant per component).
- **Precondition 2.** All communications in the pattern should exist in the application.

C. Phase 1 (Preparation)

The purpose of this initial phase is to distill essential artifacts from a security pattern for use in the subsequent phase, specifically: *pattern representation artifact*, *properties (postconditions)*, *constraints (preconditions)*. This phase serves as the preparatory step for the pattern integration process, where the critical inputs required for the process are identified and readied, effectively setting the stage for the successful application of the security pattern.

D. Phase 2 (Elicitation)

This phase aims to identify where and how the pattern will be applied. This Elicitation phase involves two key activities: *identification* and *role binding*. Initially, the architect pinpoints the software components within the application diagram that are suitable for the application of the pattern representation artifact. This step involves identifying components that are candidates to fulfill roles within the pattern representation artifacts. Following the identification of these components, the architect assigns them to their respective *roles* within the patterns, as depicted in Fig. 1 through the use of dashed arrows marked with *plays* stereotypes.

E. Phase 3 (Context Validation)

The objective of this phase is to ascertain whether the application satisfies the specified preconditions of the pattern. This phase takes as input the *Casting Diagram* and pattern *Preconditions* and provides as output a Result Artifact (that contains, for each constraint, the result of the checking). It is imperative that all preconditions are met before proceeding to the subsequent phase of the integration process, ensuring that the foundational requirements for security pattern integration are thoroughly validated.

F. Phase 4 (Merge)

The aim of this phase is to correctly integrate the pattern using the *Casting Diagram*. The Merge phase involves the integration of *Security Mechanisms*, ports, interfaces, and communication features from the Pattern Representation Artifact into the application's architectural diagram. This integration culminates in a revised version of the application, which is considered a candidate for enhanced security. This phase is crucial for ensuring that the security improvements are accurately and effectively incorporated into the application's design, aiming to elevate the overall security posture of the application. Let A be an application, P a pattern representation artifact and C a casting diagram containing a set of bindings bi . The resulting application obtained by integration process RA is defined by the algorithm in Listing 1.

```

1 Algorithm Merge
2
3 Input: A, P, C.
4 Output: RA.
5
6 RA:= duplicate(A)
7 C':= duplicate(C)
8 for each bi in C'
9   component1 = bi.component
10  patternParticipant1 = bi.patternParticipant
11
12 for each pPort in patternParticipant1
13   if pPort.communication.prts.component->includes(
14     SecurityMechanism)
15     Prt1 = RA.duplicatePort(pPort,component1)
16     securityMechanism = RA.duplicateComponent(Port.
17       securityMechanism)
18     RA.CreateCommunication(Prt1 ,securityMechanism.
19       port)
20   else
21     patternParticipant2 = Port.communication.ports.component
22     ->select(PP | PP != patternParticipant1)
23     component2 = patternParticipant.binding.component
24     Prt1 = RA.ports->select(prt1 prt. owner = component1 and
25       Prt1.communication.connects(component1 ,component2)
26     )
27   for each pOperation in Port.interface.operations
28     if Prt1.interface.operations->includes(pOperation) ==
29       false
30     RA.addOperation(pOperation , Prt1.interface)
31   endif
32   endfor
33 endif
34 endfor
35 endfor

```

Listing 1. Merge algorithm

First, the application diagram A and casting diagram C are duplicated and named RA and C' respectively. The set of bindings in C' are parsed. For each binding bi , $patternParticipant1$ and $component1$ are the pattern participant and the application component bound by bi respectively. In addition, ports owned by $patternParticipant1$ are looked up. Afterward, security mechanisms are deployed. For each port $pPort$, if $pPort$ connects $patternParticipant1$ to a security mechanism, then it is duplicated, added to $component1$ in RA , and named $Prt1$. The security mechanism is duplicated, added to RA , and named $securityMechanism$. A communication is created between $Prt1$ and a $securityMechanism$ port. Finally, the necessary operations are added to the interfaces of the application components in order to correctly call the operations of the security mechanisms. If $pPort$ does not connect $patternParticipant1$ to a security mechanism, then it is connected to another pattern participant that we name $patternParticipant2$. We name $component2$ the application component bound to $patternParticipant2$. According to the assumptions, there must be communication between application components (this is a precondition). In this case, $component1$ is connected to $component2$ via port $Prt1$. The operations of $pPort$ interface are looked up. For each operation $pOperation$ in $pPort$ interface, if $pOperation$ does not exist in $Prt1$ interface, $pOperation$ is added to the operations of $Prt1$ interface. After the merge phase, the resulting application verifies the *pattern postconditions*.

G. Phase 5 (Verification & Validation)

At this stage, the new application verifies the post-conditions. At each change (e.g., ad-hoc tailoring or the integration of another pattern), the application is validated against the postconditions. A dedicated checking module is responsible for ensuring the application adheres to these postconditions, which reflect the essential security properties of the pattern. Modifications to the application, whether due to specific customizations or the addition of another security pattern, are scrutinized under this process to ensure the security enhancements are properly maintained.

The application's compliance with pattern postconditions is assessed by ensuring: (1) the inclusion of necessary security mechanisms, like encryption, to uphold the pattern's security attributes within the application, and (2) the correct application of these mechanisms, such as encrypting messages, to safeguard data integrity and confidentiality during transmission. This process confirms both the presence and proper implementation of security measures as stipulated by the security pattern.

IV. MDE FRAMEWORK

MDE is used to support the aforementioned approach, focusing on software architecture and security patterns. For architectural modeling, we used a Domain Modeling Specific Language (DSML) based on UML modeling language

to describe software architecture using the component-port-connector fashion. For specifying and analyzing security properties, OCL is utilized, allowing for precise definition and verification of security attributes within the system.

Modeling the architecture.: In the context of Component-Based Development (CBD), the "ComponentUML" UML profile was developed to facilitate application modeling. This necessity for a specialized profile emerged during the formalization of OCL expressions, where the application of standard UML led to complexities due to the inclusion of irrelevant concepts. To streamline this process, the "ComponentUML" profile was crafted, focusing on simplification. It is built upon key UML constructs, specifically StructuredClassifiers, Messages, and Deployments, to tailor the modeling experience to CBD's specific needs.

Working example: metamodel instantiation: Fig. 2 shows the software architecture of a three-tier web application. The structure of the architecture is delineated into three fundamental types of components: *Page*, *Webapp*, and *Database*. Each of these components is interconnected through specific ports, interfaces, data types, and messages. For example, a *Page*-type component, referred to as *Webapp*, leverages a "Port Client Server" for communication with the *Webapp* component, which is classified as a *Webapp*-type. In response, the *Webapp* component utilizes a "Port Server Client" for reciprocal communication. Annotations in blue delineate various messages: "m1" describes the request dispatched to the application, "m2" illustrates the application's response, "m3" signifies the request made to the database, and "m4" represents the response received from the database. From a deployment standpoint, the architecture is supported by three nodes, which serve distinct functions: the *Browser* node, which hosts the *webpage*; the *Server* node, which is accessible via the Internet and hosts the *Webapp*; and a backend node dedicated to database hosting.

Modeling security patterns: We developed a UML profile called *SepmUML* using UML notations. *SepmUML* contains the necessary stereotypes for modeling a security pattern in UML environments. The solution of the security pattern is modeled using ComponentUML. In addition pattern integration-related concepts are specified. The specification of the UML profile and the pattern specification is out of scope in this paper and is detailed in [19]. In this process, different stakeholders interact in order to specify the pattern and its constituents including the pattern solution and the OCL constraints. Five main security mechanism categories of *SepmUML* are considered: Authenticity, Authorisation, Authentication, Cryptography, Monitoring, and Filtering. Derived from security requirements, a customized security pattern solution can be built up from a combination of these security mechanisms categories.

Working example : SSL pattern: Fig. 3 shows the solution of the SSL (Secure Sockets Layer) pattern. It has two pattern participants *clientParticipant* and *serverParticipant*. It contains the following security mechanisms. The *protocol con-*

troller is the main component that manages the other security features to facilitate the execution of the *SSL Handshake* and *SSL record* protocols. *The authenticator* validates the identity of either the client or server by verifying their certificates. *The key exchange* component calculates the key exchange for the client, which is essential for encrypting communications. While the *encryptor* applies encryption to outgoing messages utilizing a specific key and the *decryptor* utilizes a key to decrypt incoming messages. *The Signer* generates a digital signature for each message to ensure its authenticity and integrity, which accompanies the transmitted message. *The verifier* confirms the authenticity and integrity of a message by examining the digital signature that comes with it. The description of the properties is encapsulated within a comment annotated by the stereotype *PropertySpecification*, encompassing the attributes *confidentiality*, *integrity*, and *authenticity*. As shown at the bottom of Fig. 3, the pattern provides the following properties: confidentiality of messages m_1 and m_2 (**Property 1**), integrity of messages m_1 and m_2 (**Property 2**), authenticity of messages m_1 and m_2 (**Property 3**). The preconditions are: All pattern participants should be bound (one participant per component) (**Precondition 1**). Additionally, all pattern-defined communications must be present within the application, identified as (**Precondition 2**).

Casting diagrams and preconditions: Fig. 4 shows the following bindings between the SSL patterns and components of the web application where "webpage" and "webapp" play role "clientParticipant" and "serverParticipant" respectively. The casting diagram is validated against *Precondition 1* and *2*. The validation of the castings against the preconditions is done using OCL invariants. Listing. 2 and Listing. 3 validate *Precondition 1* and *2* respectively. The two preconditions have been already explained previously and recalled as comments at the beginning of the listings. In this case, the preconditions are valid so we move to the next phase.

```

1 //Precondition 1 : All pattern participants should be bound
  (one participant per component)
2 Context Castings
3 self.play->select(p1,p2 |
4
5 (p1.participant = p2.participant
6 implies
7 p1.component = p2.participant)
8
9 and
10
11 (p1.component = p2.participant
12 implies
13 p1.participant = p2.participant)
14
15 and
16 p1.participant.structureContainer.participants->forAll(
  participant | self.play->exists(p_ | p_.participant =
  participant)))

```

Listing 2. Precondition 1: All pattern participants should be bound (one participant per component)

```

1 //Precondition 2: all communications in the pattern should
2 exist in the application

```

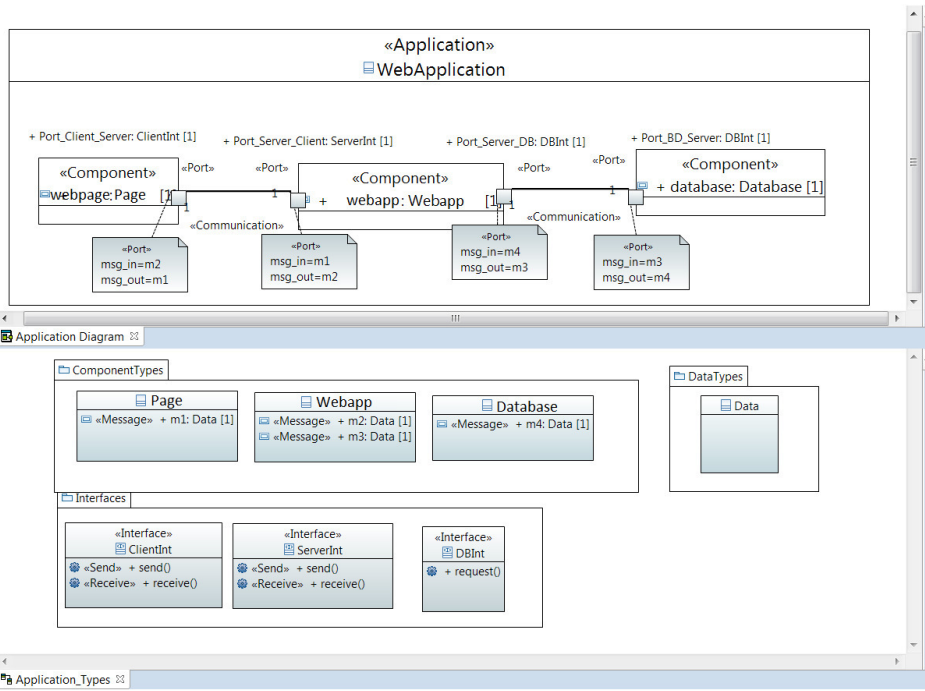


Fig. 2. Architecture Model and Component Types of a Web Application [18]

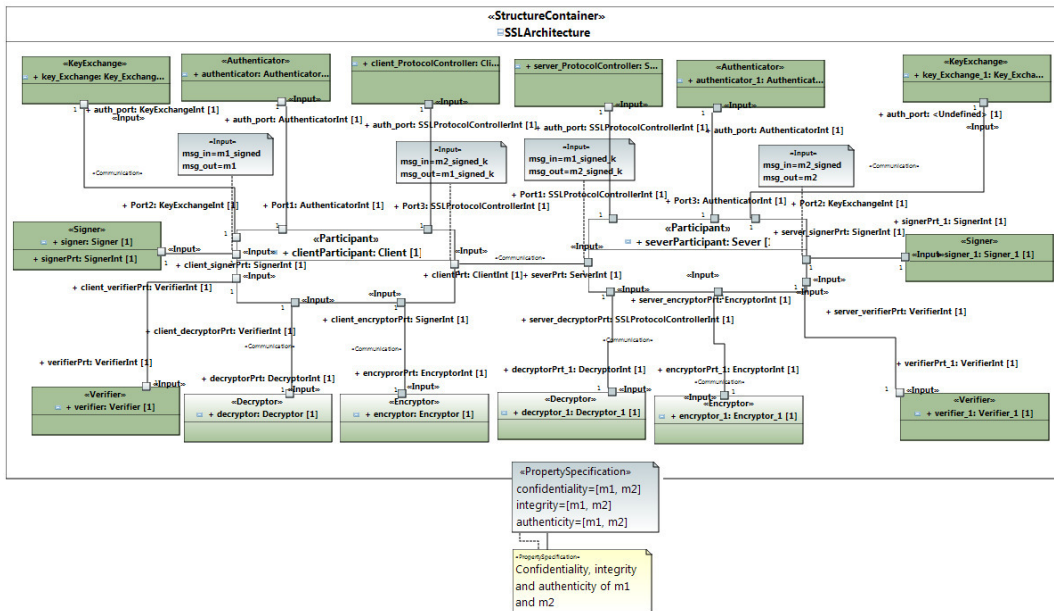


Fig. 3. SSL Pattern Solution




Fig. 4. Casting Diagram for Application Communication Using SSL Security Pattern

```

3 Context Castings
4 self.play->forAll(p1,p2 |
5 (p1.'<>'(p2) and p1.participant.ports->exists(p1_in | p2.
  participant.ports->exists( p2_in | p1_in.communication
  = p2_in.communication)))
6 implies
7 p1.component.ports->exists(p1_in | p2.component.ports->
  exists( p2_in | p1_in.communication = p2_in.
  communication))
8 )
  
```

Listing 3. Precondition 2: all communications in the pattern should exist in the application

Merge: In this phase, a transformation from Model to Model is executed, utilizing the procedure outlined in Listing 1 written with QVT (Query, View, and Transformation)². This transformation process inputs the application depicted in Fig. 2 along with the castings shown in Fig. 4. The result of the Merge is a refreshed application model illustrated in Fig. 5. Furthermore, the initial application's types and interfaces undergo modifications with the addition of new attributes and operations, highlighted in green. Additionally, new interfaces pertinent to security mechanisms have been introduced.

Context checking: At this phase, the new application undergoes a verification process to ensure it meets the postconditions. Whenever a modification occurs, such as ad-hoc adjustments or the amalgamation of an additional pattern, the application's conformity with the postconditions is reassessed. This evaluation employs the OCL invariant as detailed in Listing. 4. If the postconditions still hold, the change is accepted and committed. Else, the change is dismissed. The validation process involves the application being checked against *Property 1, 2, and 3* through the confirmation of the following criteria:

- There must be one *encryptor* and *signer* mechanisms for each component sending messages m_1 and m_2 .
- Messages m_1 and m_2 must be encrypted before being sent.

```

// Postconditions validation
Context Application
self.base_class.ownedAttributes->select(c | c.isOclAsKind(
  Comment))->select(c |
c.isStereotypeApplied( PropertySpecification ))->forAll( p |
p.confidentiality->forAll( m_confidential |
  
```

²<https://www.omg.org/spec/QVT>

```

// There exists a component that produces m_confidential
  and an encryptor connected to this component that
  encrypts m_confidential
self.components->exists( c1 |
c1.ports.msg_out = mconfidential
and
self.components->exists( enc1
enc1.isOclKindOf(Encryptor)
and
c1.ports->exists(c1_in | enc.ports->exists(enc_in | c1_in.
  communication =
  enc_in.communication )))
and
p.integrity->forAll( m_integrity |
// There exists a component that produces m_integrity and a
  signer connected to this component that signs
  m_integrity
self.components->exists( c1 |
c1.ports.msg_out = m_integrity
and
self.components->exists( signer |
signer.isOclKindOf(Signer)
and
c1.ports->exists(c1_in | signer.ports->exists(signer_in |
  c1_in.communication =
  enc_in.communication )))
  
```

Listing 4. OCL Queries for pattern property verification

V. TOOL SUPPORT

The proposed toolchain is designed to support the proposed metamodellers (ComponentUML and SepmUML) and Model-To-Model transformations. To support our approach, tools must fulfill the following key requirements:

- Allow the creation of a custom UML profile and UML models.
- Support the implementation of a repository to store pattern models and the related model libraries for classification and relationships.
- Support the access to the repository. Create views on the repository according to its APIs, its organization, and the needs of the targeted system engineering process.
- Enable transformations of the pattern models from the repository format into the target-modeling environment.
- Enable the creation of System of Patterns configuration models in the target-modeling environment.
- Enable the integration of patterns imported from the model repository into application models.

In our case, the following support tools have been chosen:

- UML modeling environment: Papyrus³ (Existing)
- Model Repository : SEMCOMDT⁴ (SEMCO Model Development Tools, IRIT's editor and platform plugins) is used to support pattern repository (Existing).
- Selection, Instantiation, and Integration of Pattern Models: Semco4Papyrus (Implemented)

VI. CASE STUDY: SCADA SYSTEM

This section evaluates the applicability and efficacy of our proposed methods through the study of a SCADA system. SCADA systems diverge significantly from conventional IT

³<https://eclipse.org/papyrus/>

⁴<http://www.semcomdt.org>




Fig. 5. New Application diagram

frameworks, including web applications, showcasing particular and rigorous security demands. The purpose of this evaluation is to illustrate the capability of our suggested approaches in meeting the sophisticated security challenges that are fundamental to these vital infrastructure components.

Description and modeling: In our study, we explore a simplified SCADA system tailored for smart grid applications. This system is designed to perform essential functions such as: (1) Perform control, (2) Poll Data, (3) System Start-up/shutdown, (4) Adjust Parameter Settings, (5) Log Field Data, (6) Archive Data, (7) Trigger Alarm, (8) Perform Trending (e.g., Select Parameters, Display Parameters, Zooming and Scrolling). Fig. 6 depicts the software architecture model.

A set of patterns is selected from the model repository and then instantiated in the modeling environment Papyrus⁵ (see Fig. 7). The security patterns are:

- *Secure Communication* that ensures that data passing across a network is secure. It can be refined by two alternative patterns: SSL and IPSec.
- *Firewall* that restricts access to the internal network. It can be refined by the following alternative patterns: Packet Firewall and Stateful Firewall.
- *Intrusion Detection System (IDS)* in order to stop malicious payloads.
- *Authorization* that ensures that only authorized users are allowed access.
- *RBAC (Role-Based Access Control)* that allows the creation of a set of roles. Each role has a set of defined rights.
- *Logger and Auditor.* that allows the logging of actions.

Amongst the different Patterns, two have been selected: SSL and Packet Firewall patterns and integrated into the SCADA system software architecture. The integration aims at protecting: (1) the communication between the SCADA server and PLCs (Programmable Logic Controller) against information disclosure and spoofing (e.g., Man-In-The middle attacks), (2) the SCADA server and PLCs against Denial of Service attacks.

Assessment: For the assessment, we use MBTA [18] which is a framework for detecting threats based on OCL. The aim is to analyze the software architecture before and after pattern integration. The approach and threat categories are detailed in [18]. Fig. 8 presents a comparative analysis of the frequency of threats categorized by type, before and following the adoption of security patterns. The implementation of the SSL pattern eradicates threats associated with Man-In-The-Middle and Tampering (during transmission) by ensuring the confidentiality and integrity of communications between the SCADA server and the PLCs. However, vulnerabilities to Denial of Service and Injection attacks persist, affecting four software components: HMI (Human-Machine Interface), Trending, LogDisplay, and AlarmDisplay; due to their exposed public ports and the absence of Firewall and Authorization safeguards.

VII. CONCLUSION

In this paper, we have proposed an approach and tool support for integrating proper security patterns into software architectures, aligning with OMG (Object Management Group) standards such as UML, profile extension mechanism, and OCL. Our process utilizes merging and OCL verification strategies within an MDE based approach to facilitate the integration of security patterns. The challenge of pattern integration lies in seamlessly embedding all components of a pattern into an application, ensuring system integrity and quality are maintained, and affirming the enhancement of the system with the new properties introduced by the pattern. While the majority of research in this area concentrates on utilizing merging techniques for the incorporation of patterns into applications, our study emphasizes the Verification & Validation phase to validate the integration. Through the examination of a case study and the conducted assessment, we have pinpointed limitations in the current iteration of MBSPI. Our future directions include the proposition of pattern bindings to users, tailored according to the component types targeted for pattern application. During the integration phase, our focus has been predominantly on the structural aspects of architecture, as

⁵<https://eclipse.org/papyrus/>

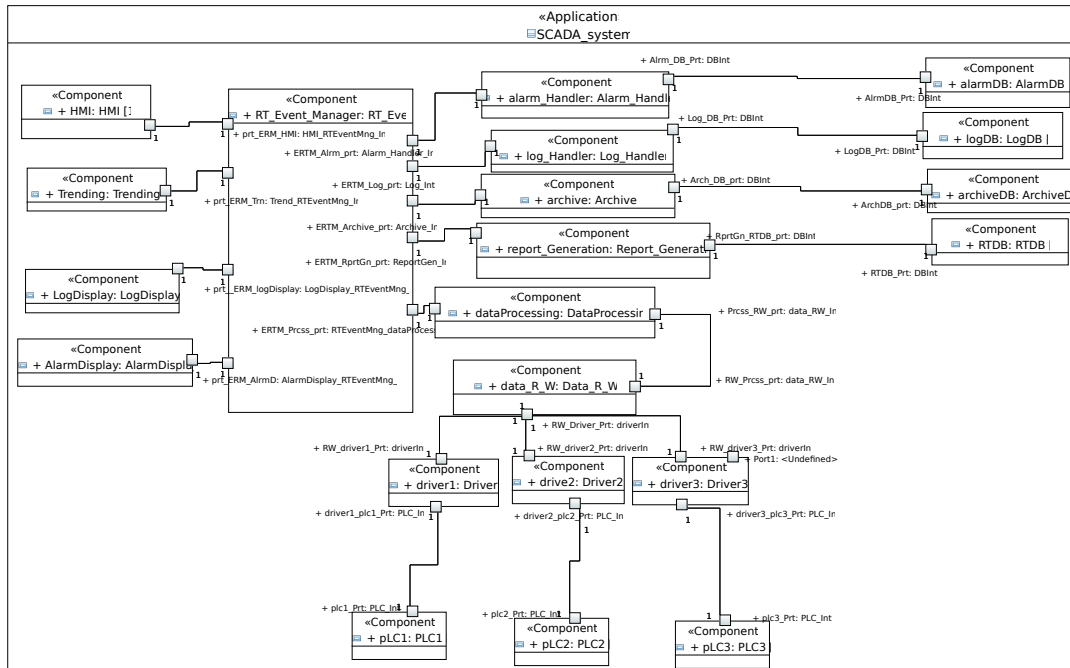


Fig. 6. SCADA software architecture model

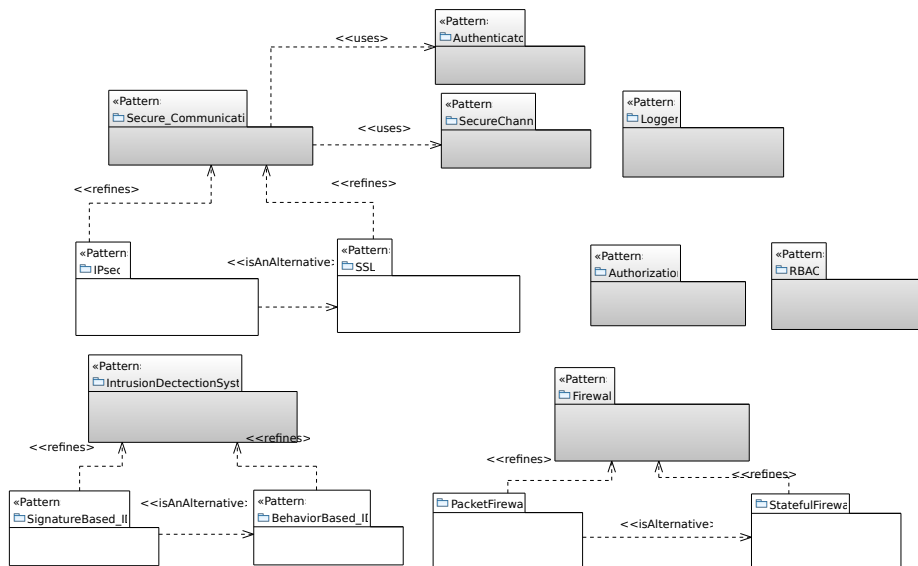


Fig. 7. Selected Security Patterns

depicted in UML composite diagrams. Moving forward, it is imperative to also account for behavioral considerations during the integration process. Currently, messages are treated as structural elements within our solution; however, to adequately represent the sequencing and interaction of messages, we propose the utilization of dedicated diagrams, such as the UML sequence diagram, which better captures message order and interaction dynamics.

REFERENCES

- [1] E. Fernandez-Buglioni, *Security patterns in practice: designing secure architectures using software patterns*. John Wiley & Sons, 2013.
- [2] H. Washizaki, T. Xia, N. Kamata, Y. Fukazawa, H. Kanuka, T. Kato, M. Yoshino, T. Okubo, S. Ogata, H. Kaiya, et al., "Systematic literature review of security pattern research.," *Information*, vol. 12, no. 1, pp. 2078–2489, 2021.
- [3] D. Manolescu, W. Kozaczynski, A. Miller, and J. Hogg, "The Growing Divide in the Patterns World," *IEEE Software*, vol. 24, pp. 61–67, July 2007.
- [4] D. Serrano, A. Mana, and A.-D. Sotirious, "Towards Precise and Certi-




Fig. 8. Threats Counts Before and After Security Pattern Integration

fied Security Patterns,” in *Proceedings of 2nd International Workshop on Secure systems methodologies using patterns (Spattern 2008)*, pp. 287–291, IEEE Computer Society, September 2008.

- [5] J. F. Ruíz, M. Arjona, A. Maña, and N. Carstens, “Secure engineering and modelling of a metering devices system,” in *2013 International Conference on Availability, Reliability and Security, SecSE’13*, pp. 418–427, IEEE, 2013.
- [6] A. Maña, E. Damiani, S. Gürgens, and G. Spanoudakis, “Extensions to Pattern Formats for Cyber Physical Systems,” in *Proceedings of the 31st Conference on Pattern Languages of Programs*, no. 15 in PLOP’14, pp. 15:1–15:8, ACM, 2014.
- [7] P. H. Nguyen, K. Yskout, T. Heyman, J. Klein, R. Scandariato, and Y. L. Traon, “SoSPa: A system of Security design Patterns for systematically engineering secure systems,” in *2015 ACM/IEEE 18th International Conference on Model Driven Engineering Languages and Systems (MODELS)*, pp. 246–255, Sept. 2015.
- [8] D. Mouheb, C. Talhi, M. Nouh, V. Lima, M. Debbabi, L. Wang, and M. Pourzandi, “Aspect-Oriented Modeling for Representing and Integrating Security Concerns in UML,” in *Software Engineering Research, Management and Applications*, no. 296 in Studies in Computational Intelligence, pp. 197–213, Springer Berlin Heidelberg, 2010.
- [9] J. M. Horcas, M. Pinto, and L. Fuentes, “An Aspect-Oriented Model transformation to weave security using CVL,” in *2014 2nd International Conference on Model-Driven Engineering and Software Development (MODELSWARD)*, pp. 138–150, Jan. 2014.
- [10] G. Georg, I. Ray, K. Anastasakis, B. Bordbar, M. Toahchoodee, and S. H. Houmb, “An aspect-oriented methodology for designing secure applications,” *Information and Software Technology*, vol. 51, pp. 846–864, May 2009.
- [11] S. Peldszus, “Model-driven development of evolving secure software systems,” in *Combined Proceedings of the Workshops at Software Engineering 2020 Co-located with the German Software Engineering Conference 2020 (SE 2020) (R. Hebig and R. Heinrich, eds.)*.
- [12] X. Zheng, D. Liu, H. Zhu, and I. Bayley, “Pattern-based approach to modelling and verifying system security,” in *15th IEEE International Conference on Service Oriented Systems Engineering (SOSE)*, pp. 92–102, 2020.
- [13] H. A. Alhamad and M. M. Hassan, “Aspect-oriented models-based framework to secure intelligent systems,” in *Proceedings of the 8th International Conference on Computer Technology Applications (ICCTA)*, pp. 249–262, 2022.
- [14] A. Armouh, “Towards the integration of security and safety patterns in the design of safety-critical embedded systems,” in *4th International Conference on Applied Automation and Industrial Diagnostics (ICAAID)*, vol. 1, pp. 1–6, 2022.
- [15] B. Hamid, C. Percebois, and D. Gouteux, “A Methodology for Integration of Patterns with Validation Purpose,” in *Proceedings of the 17th European Conference on Pattern Languages of Programs (EuroPLOP)*, pp. 1–14, ACM, 2012.
- [16] R. Abdallah, A. Motii, N. Yakymets, and A. Lanassee, “Using model driven engineering to support multi-paradigms security analysis,” in *Model-Driven Engineering and Software Development: Third International Conference, MODELSWARD 2015, Angers, France, February 9-11, 2015, Revised Selected Papers 3*, pp. 278–292, Springer, 2015.
- [17] A. Motii, B. Hamid, A. Lanassee, and J.-M. Bruel, “Towards the integration of security patterns in UML component-based applications,” in *Joint Proceedings of the Second International Workshop on Patterns in Model Engineering and the Fifth International Workshop on the Verification of Model Transformation*, vol. 1693 of PAME ’16, pp. 2–6, CEUR-WS.org, 2016.
- [18] A. Motii, “Mbta: A model-based threat analysis approach for software architectures,” in *42nd International Conference on Computer Safety, Reliability, and Security (SafeComp)*, pp. 121–134, 2023.
- [19] A. Motii, B. Hamid, A. Lanassee, and J. M. Bruel, “Guiding the selection of security patterns for real-time systems,” in *21st International Conference on Engineering of Complex Computer Systems (ICECCS)*, pp. 155–164, 2016.

An Ontology to Understand Programming Cocktails

Alvaro Costa Neto
0000-0003-1861-3545

Research Centre in Digitalization and
Intelligent Robotics (CeDRI)
Laboratório para a Sustentabilidade e Tecnologia
em Regiões de Montanha (SusTEC)
Instituto Politécnico de Bragança
Campus de Santa Apolónia, 5300-253 Bragança, Portugal

ALGORITMI Research Centre / LASI, DI
University of Minho, Braga, Portugal

Instituto Federal de Educação,
Ciência e Tecnologia de São Paulo
Barretos, Brazil

Email: alvaro@ifsp.edu.br

Maria João Varanda Pereira
0000-0001-6323-0071

Research Centre in Digitalization and Intelligent Robotics (CeDRI)
Laboratório para a Sustentabilidade e Tecnologia
em Regiões de Montanha (SusTEC)
Instituto Politécnico de Bragança
Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Email: mjoao@ipb.pt

Pedro Rangel Henriques
0000-0002-3208-0207

ALGORITMI Research Centre / LASI, DI
University of Minho, Braga, Portugal
Email: prh@di.uminho.pt

Abstract—An ever-growing landscape of programming technologies (tools, languages, libraries and frameworks) has rapidly become the norm in many domains of computer programming—Web Development being the most noticeable example. The concurrent use of many compartmentalised technologies has advantages: it allows for flexibility in implementation, while also improving reusability. On the other hand, this proliferation tends to create convoluted development workflows that must be (painstakingly) planned, managed and maintained. The combination of multiple languages, libraries, frameworks and tools (*Ingredients*) in a single project effectively forms a *Programming Cocktail*, that can rapidly become cognitive and financially onerous. Aiming at understanding these complex situations, an ontology was created to provide a formal and structured analysis of these cocktails. It emerged from a survey of technologies that several companies are currently using to develop their systems, and aims to provide support for better understanding, classifying and characterising *Programming Cocktails*. This paper presents not only the ontology itself, but also the consequent knowledge that was constructed and structured through its development.

Index Terms—Ontology, Programming Cocktails, Software Development, Programming Technologies, Knowledge Construction

I. INTRODUCTION

THE DEVELOPMENT process of an application invariably requires the use of certain technologies, such as languages (for programming, specification *etc.*), libraries, frameworks and tools. This process may be monotonic, requiring no more than a base language and, occasionally, a handful of libraries. On the other hand, it may also be plural and

This work has been supported by FCT – Fundação para a Ciência e Tecnologia within the R&D Units Project Scope: UIDB/00319/2020.

The work of Maria João and Alvaro was supported by national funds through FCT/MCTES (PIDDAC): CeDRI, UIDB/05757/2020 (DOI: 10.54499/UIDB/05757/2020) and UIDP/05757/2020 (DOI: 10.54499/UIDP/05757/2020); SusTEC, LA/P/0007/2020 (DOI: 10.54499/LA/P/0007/2020).

polyglot, with several frameworks, libraries, tools and components (modern Web applications being the most prominent examples).

As is typical in the latter case, whenever a project demands—or is propelled by—the presence of multiple technologies, those involved in the construction of the application must learn, use and manage them. The epistemological challenges that arise in these situations resemble the ones that have been studied in Computer Programming Education for decades [1]. These studies range from tools to aid students and teachers [2]–[5], educational methodologies [6]–[8], success and failure factors [9], [10] to more psychological endeavours [11]–[14]. Technologies that are unfamiliar to programmers must be learnt and understood [9], [15] during the entire life cycle of an application. Either in the initial development phases of a project, or when technology adoption changes, knowledge must be constructed for the learning process to happen. In the presence of several technologies, a new caveat appears: beyond understanding each one, programmers must also manage a surge in cognitive burden as their brains are required to cope with alternating mental models.

Research into dealing with these challenges usually present themselves as comparative surveys [16]–[18] that list different programming technologies and their main characteristics. They usually aim to establish a clear landscape and support decisions on which technologies are best suited to specific contexts based solely on individual properties of each technology. Inherently, these studies fail to take into account any possible combination thereof, focusing their efforts in relating and comparing pre-determined aspects. It then becomes clear that a comparative study is not enough to understand how these programming technologies relate to each other in real-life scenarios. The concepts that relate to these technologies, and their interconnections must be formally and structurally

mapped. Knowledge must be constructed to cover not only each Ingredient (technology), but the Cocktail (combination) itself. A possible answer to this challenge relies in the use of ontologies [19], [20], a formal method to structure knowledge, to conceptualize and instantiate information from Cocktails, establishing reasonable inferences on its landscape of programming technologies.

This article is divided into four more sections. Section II presents the definitions for Programming Cocktails, their Ingredients and other related concepts. Section III details a survey of real-life Cocktails conducted with several Software Companies, and the overall results that have been observed. Section IV presents the ontology that was created to formally analyse and understand Programming Cocktails, the concept of a Cocktail Identity Card, and what knowledge was constructed around and through them. Finally, Section V concludes the paper with a summary of lessons learned, and presents the next steps in the research of Programming Cocktails.

II. PROGRAMMING COCKTAILS

Before delving into the intricacies of application development and the complex relations between the components that are used to build them, it is of good measure to define what *Programming Cocktails* and *Ingredients* actually mean.

It is comprehensible that the use of such relaxed terms to describe logical and structured concepts might seem as a stretch (or even sarcastic) at first sight. Maybe just an analogy, that is furiously gripping itself on the edge of an undeniably sharp, sleek and mathematically sound cliff. Nevertheless, it is in fact very meaningful to this article's context and objectives. Anyone who has ever tried to concoct actual cocktails should be able to describe them by more than a list of components. The results are sometimes clean and homogeneous, with strong and decisive tastes. In other cases, the components barely mix together, presenting fuzzy (even chunky) separations that stubbornly remain. In the worst scenarios, when the list of ingredients, their measures, and combinations are poorly chosen, the final result may become undrinkable.

Analogously, the term *Programming Cocktail* defines a combination of computer programming technologies—such as programming languages, libraries and frameworks—that is used to develop specific software applications. *Ingredients* are the components of a *Cocktail*. It is important to note that a Cocktail is associated with a specific application or service, and its Ingredients may also appear in the Cocktail for other application under the same development context¹. Suppose a company develops three applications:

- **Application A:** HTML, CSS, JavaScript, and ReactJS;
- **Application B:** HTML, CSS, JavaScript, MySQL, and PHP;
- **Application C:** C++, and Unity.

It might seem that, as a whole, there is one Cocktail for the company: the union of the sets formed by each application's

Cocktail. Nonetheless, for the purposes of this study, each Cocktail is taken independently, even if it means to consider Ingredients more than once in the same development context. In short, there are three Programming Cocktails in the previous example, one for each application (A, B, and C).

As is expected, a few decisions had to be made while defining these terms. The first and foremost was: which development technologies should be considered Ingredients? At first sight, there are countless technologies that are involved in the development of an application. From standard and well-known programming languages, through Domain-Specific Languages (DSL) for diverse specifications, configuration and communication; to niche libraries, full-stack frameworks, editors and debuggers, the list of candidates to be identified as Ingredients is varied and long. A qualitative threshold was defined to separate what would be considered part of a Cocktail. A programming technology was identified as an Ingredient only if it is *directly applied to the development² process of an application*.

On the other hand, several technologies are commonly used during deployment or execution of an application, such as Database Management Systems (DBMS), queue coordinators, *etc.* Despite their influence on the design and implementation of an application, these technologies are not considered Ingredients, they are *Resources*. Examples include: Apache Web Server [21], ActiveMQ [22], MySQL [23], and memcached [24].

The second decision concerned the definition of categories for the Ingredients. For the purposes of this study, an Ingredient may be categorized as one of four possibilities:

- **Language:** encompasses any kind of text or graphics-based language. May be used for programming, specification, description, communication, scripting, so on and so forth. Examples: C [25], Python [26], HTML [27], CSS [28], SQL [29], *etc.*
- **Library:** a portion of code (either in source form or pre-compiled) that augments programming languages and their standard libraries with extra functionality. Examples: LibSSH [30], RayLib [31], *etc.*
- **Framework:** scaffolding augmentations to programming languages. Albeit similar to libraries, frameworks add functionality while imposing some form of structure to the source code (syntactic, semantic, or paradigmatic) or the use of pre-defined components³. Examples: SwiftUI [32], React Native [33], *etc.*
- **Tool:** specifies any tool that is directly used for development, such as editors, Integrated Development Environments (IDE), debuggers *etc.*

There might be cases in which the borders between these categories become tenuous. In these situations, an Ingredient that has multiple roles in the development process might need

²*Development* here indicates a generalised concept which includes, but is not limited to, programming tasks.

³Given that there is no standard for distinguishing between libraries and frameworks, this definition may collide with others'.

¹*Development context* represents the set of factors that influence the actual construction of an application, including, but not limited to, its team, technologies, tools, organization, and requirements

to be either sliced into its constituent parts, or included in more than one category. As an example, testing frameworks, such as JUnit [34], usually include both libraries and servers to allow for concurrent testing. These Ingredients could possibly be separated into their individual roles (*JUnitLibrary* and *JUnitServer*, per example) or included in both categories (Library and Tool).

The reasoning behind these decisions became clear through the construction of the ontology (explained in Section IV), with its foundational rationale extracted from real-world Cocktails, surveyed from several multi-national companies in Portugal.

III. COCKTAILS ASSEMBLAGE

As previously stated, obtaining the current uses of Programming Cocktails was paramount to establishing an overall picture of computer programming technologies. To this intent, several companies were contacted in a survey for information about which Programming Cocktails they have used. Their feedback allowed for the construction of the ontology's main concepts (presented in Section IV) and the consequent structuring of the knowledge surrounding Programming Cocktails.

A. Survey

Starting in October 2023, several companies that have offices in Portugal were contacted via email for a survey of Programming Cocktails. The email (shown in the Appendix) described the context of the study and asked for the programming technologies each company has used, divided by applications in which each Cocktail was used. Companies were specifically asked to answer informally via email, in order to stimulate participation and consequently obtain faster and more numerous responses. Given previous experiences, online survey questionnaires, such as those created via platforms akin to Google Forms⁴ tend to be postponed, resulting in fewer answers. While the amount of companies that responded was far from ideal, this number would possibly be even lower if a formal system was used.

Up to the time of this paper's submission⁵, 213 companies were contacted and of those, 15 responded with several Cocktails they have used in the past, or still use in the present.

A few important considerations:

- Given the informal nature of the survey, some answers had to be either supplemented (in the case of obvious missing elements, such as Cocktails with React that missed JavaScript) or followed up with further communication;
- Some answers pointed to the fact that the *borders between some systems are a bit fuzzy*, and their Cocktails represent overall divisions that are shared between groups of applications. Such is the case with systems that are heavily structured around micro-services, as example;

- Despite the fact that only Portuguese offices of the surveyed companies were contacted, the majority of them have international endeavours or are multinational themselves, which reduces the locality bias of the answers;
- Exhaustiveness was never the goal for the survey. Given that the programming technologies landscape is ever changing in a fast pace, the survey was designed to support the construction of knowledge about Programming Cocktails, which in turn, may eventually be applied to future works and studies.

Currently, 49 Programming Cocktails have been obtained, spanning a total of 124 different Ingredients and Resources, that range from programming languages and frameworks, to database management systems and resource cache management applications.

B. Data Overview

As previously mentioned, statistical data analysis is not the main goal of this study. Nonetheless, a few statistical facts can be extracted from the Programming Cocktails that were gathered in the survey.

In order to better organize the survey results, an online spreadsheet⁶ was created, listing Ingredients on the lines and Cocktails on the columns (Table I represents a summarised example of the actual spreadsheet). Column *A* contains the names of the Ingredients. The columns from *B* to *D* categorize each Ingredient into, respectively, its type⁷ (*Language*, *Library*, *Framework* or *Tool*), the Language with which it was used, and the Task it was applied to (Tasks will be further explained in Subsection IV-B). In the eventual case of an Ingredient either being used with more than one Language, or applied to more than one Task, its line would be duplicated and its categories adapted as needed. An hypothetical example would be the .NET Framework, which can be used with several different programming languages, and would require such treatment. Finally, from *E* onwards, each Cocktail was listed in its own column, with their Ingredients' rows marked to represent their inclusion. As an example, in Table I, the column *App2* represents the second Cocktail that was gathered and includes both C# and YAML.

In total, 63 different Ingredients have been collected:

- 23 Languages;
- 14 Libraries;
- 22 Frameworks;
- 4 Tools.

As per Resources, 61 have been gathered. Overall, each Cocktail has an average of 8 Ingredients and Resources, which a 4.6 standard deviation. Some other observations include:

- 2 Cocktails are based either on *Low Code* or *No Code* Ingredients;
- Most of the Cocktails belong to Web Development (32 in total);

⁴Available at: <https://www.google.com/forms/>

⁵All data discussed in this paper should be considered from the same time period, unless stated otherwise.

⁶A read-only version is available at: <https://bit.ly/4aFjSjj>

⁷Resources have also been included in the spreadsheet and are categorized as such, despite not being Ingredients *per se*.

TABLE I
SUMMARISED EXAMPLE OF THE COCKTAILS SPREADSHEET.

Ingredient	Type	Language	Task	App1	App2	...	AppN
.NET	Framework	C#	Full-stack development			...	X
C#	Language	C#	Server implementation		X	...	
C#	Language	C#	Full-stack development	X		...	
			⋮				
YAML	Language	YAML	Communication		X	...	

- The most frequent Ingredients are (from most to least used):
 - **Language:** HTML, CSS, JavaScript, SQL, C#;
 - **Library:** OData, Bootstrap, (all the other tied in one use);
 - **Framework:** React, Node.js, .NET, Angular, ASP.NET;
 - **Tool:** Visual Studio, Visual Studio Code, PowerPages, Liferay;
- The most frequent combination of a programming language and a framework is JavaScript with React, followed by JavaScript with Node.js, C# with .NET, and JavaScript with Angular;
- 7 Cocktails use only one language;
- At the time of writing, no Cocktails have been collected that directly apply any Artificial Intelligence (AI) support or technology.

As previously mentioned, micro-services architectures presented a challenge in defining borders between systems—and consequently, their Cocktails. In these cases (5 in total), their Cocktails were defined taking into account a group of micro-services that implement logical parts of the whole system. The logic behind this definition was dependent on the system itself, and as such, stipulated by the company that provided the Cocktails.

IV. ONTOLOGY FOR PROGRAMMING COCKTAILS

The data collected through the survey has a purpose: to allow for better understanding of Programming Cocktails, which entails the construction and structuring of knowledge. Our research group has had several interactions with and has made several contributions to the study of ontologies [35]–[41], both in their construction and definition. Consequently, from several approaches that could be applied to achieve the construction of knowledge about Programming Cocktails, an ontology seemed a straightforward and appropriate choice. It allowed for the formal definition of Programming Cocktails’ main concepts, the generation of *Identity Cards*, an ontology-based characterisation mechanism for Cocktails, and the organization of their Ingredients.

Moreover, the ontology will be paramount for future use in coming studies, that will deal with the evaluation of Cocktails in cognitive load metrics. The construction of the ontology was then, in practice, a two-fold endeavour, as it aided in structuring and understanding the data that was collected

through the survey (its initial goal), while also providing a foundation on which several studies might surge.

A. *OntoDL*

The initial version of the ontology was created using a spreadsheet to organize and list its concepts, relations, instances, and connections. While suitable for the beginning phases, when the number of elements was small, as the ontology grew it became evident that other solutions would offer better scaling and future-proofing. A visual representation would be ideal to quickly present the connections between the ontology’s elements. Given previous experiences and its simple yet capable syntax, *OntoDL* [42] was chosen as the main source for the definition and instantiation of the ontology. *OntoDL* is a Domain-Specific Language (DSL) that was created for modelling ontologies, as an alternative to more verbose options such as the Web Ontology Language (OWL) [43]. It has been used in several projects, including the *WebOntoDL* application⁸, which can interpret ontologies written in *OntoDL* and translate them to several other formats, such as *DOT*⁹ and *OWL*. It contains a syntax that is reminiscent of the mathematical formal definition of ontologies. It also allows the use of keywords in Portuguese or in English, which could be a beneficial factor for exchange and contribution from third parties.

The basis for *OntoDL*’s syntax relies on five main structures: the name of the ontology, the list of concepts, the list of individuals, the list of relation types and the triples that actually declare relations. Listing 1 shows the basic declaration for each structure. A few basic rules:

- The language follows basic principles of ignoring whitespaces and line-breaks, as well as the use of curly brackets as group delimiters;
- The order of the declarations matters;
- Concepts, relationships and triples are mandatory;
- Comments are line based, beginning with the percent sign (%) and ending with a line-break;
- There are pre-defined relation types for specialization (*isa*), composition (*prof*), and instantiation (*iof*);
- Triples are directional, and may be formed using any combination of concepts and individuals.

⁸Available at: <https://webontodl.epl.di.uminho.pt>

⁹*DOT* is a graphics format written in plain text that is used to define visual elements in a diagrammatic form. It is part of the *Graphviz* project, available at <https://graphviz.org>

Listing 1 Basic syntax for OntoDL.

```

% Identifiers must follow the C standard.
Ontology OntologyName

% The order of declarations matters.
concepts { Concept1, Concept2, ... }

individuals { Individual1, Individual2, ... }

% Whitespace and line-breaks are ignored.
relationships {
  RelationType1,
  RelationType2,
  ...
}

triples {
  % Triples may relate both concepts and
  individuals, in any combination.
  Concept1 =RelationType1=> Concept2;
  Individual4 =RelationType2=> Individual5;
  % Specialization.
  ConceptChild =isa=> ConceptFather;
  % Composition.
  ConceptPart =pof=> ConceptWhole;
  % Instantiation.
  IndividualA =iof=> ConceptA;
}
% The period indicates the end of the ontology.
.

```

There are more rules for defining properties, axioms and other elements, that have not been used in this paper. With the foundation firmly established on OntoDL, the first step to create a valid ontology that would allow for reasoning on the surveyed Programming Cocktails was to model its main concepts.

B. Open Conceptual Model

Before delving into the actual ontology and its concepts, it is important to establish a graphical notation that will be used and referenced in figures that represent them. Fig. 1 shows the basic elements and how they are graphically styled, more specifically for OntoDL's pre-defined relations. Other relationships use a simple arrow. This notation is emblematic because it allows for quick identification of element types (concepts or individuals), and pre-defined relation types that carry specific semantics (instantiation, specialization and composition).

The main concepts of the ontology were incrementally created. The first concepts meant to establish a foundation, based on the fact that each *Cocktail* is directly associated to a *System* that either is or has been under *Development*. Listing 2 and Fig. 2 present them.

The next step in the development of the conceptual model was the addition of the Ingredients and their types. The results of the survey reaffirmed the initial proposal for their types (Language, Library, Framework and Tool), but also highlighted a basic problem: some tools that were listed by the companies did not participate directly in the development process. In some cases, such as the main Operating System that programmers chose or a note taking application, it was evident

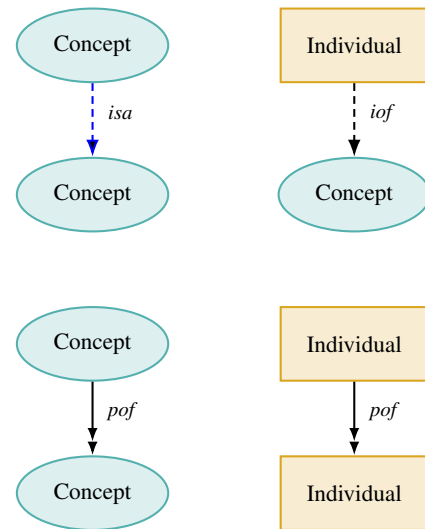


Fig. 1. Graphical notation for the ontology.

Listing 2 Initial concepts for the ontology.

```

Ontology Cocktails

concepts { System, Development, Cocktail }

relationships { uses, requires }

triples {
  % Foundation.
  System =requires=> Development;
  Development =uses=> Cocktail;
}
.

```




Fig. 2. Initial concepts for the ontology.

that their inclusion as Ingredients would be a stretch that could backfire in later developments. Other elements were not so easy to distinguish, such as Database Management Systems, and Queue Managers.

At that moment, a decision had to be made: would these tools be included as Ingredients? Which category would they belong to? At first, the intention was to include all of them as Tools, but as more Cocktails were obtained, it became clear that this choice could easily distort the category's meaning. In the end, in order to avoid this negative effect, the concept of a *Resource* was created to implement this solution. A Resource represents an external system (or service) that is used at runtime by the application, but that does not participate directly during the development process. For obvious cases, such as Operating Systems, the separation between a Tool and a Resource was clear. Alas, that was not always the case. Some Resources required some kind of implementation in the development phases, such as a communication library, or a configuration language. In these cases, the runtime of the system (or service) was considered a Resource, while any mandatory Application Programming Interface (API), library, framework or language that was used to interact with it was considered an Ingredient (of the correct type). As an example, in order to communicate with MySQL, an external library or framework (such as *libmysqlclient*) is usually included as part of the project. Given that the runtime of the MySQL server provides support for the *execution* of the application, it would be considered a *Resource*. On the other hand, *libmysqlclient* would be considered a *Library*, as it participates directly in the development phase to program the interaction to the server. In order to highlight this difference in purpose (supporting the application execution *versus* the development process), the concept of the Resource was moved from its initial relationship (a specialization of Ingredient) to a supporting role to the System itself. Listing 3 and Fig. 3 show the inclusion of these concepts to the ontology. It can be seen that while the concepts of Language, Library, Framework and Tool are specializations of the more general concept of an Ingredient, Resource is directly connected to the System concept.

The last two additions to the conceptual model consisted in a series of relationships that highlighted the central role of Languages in the Cocktail, and the definition of Tasks. Languages are usually the central element in the development of almost any kind of System. In fact, it is very unusual that a Language choice will depend on other types of Ingredients, such as Frameworks, or Libraries. The reverse, although, is commonplace: the choice of a Language usually dictates which other Ingredients will be part of the Cocktail.

In order to illustrate and define Language's central role, three relationships were added, each connecting one of the other Ingredient types to it (see the *extends*, *encloses* and *supports* relationships that terminate in *Language* in Listing 4 and Fig. 4).

The definition of Tasks and how they were modelled was the last step in the construction of the conceptual model for the ontology. As with any specification for concepts that rely

Listing 3 The inclusion of Ingredients and Resources to the ontology.

```

Ontology Cocktails
concepts {
  System,
  Development,
  Cocktail,
  Resource,
  Ingredient,
  Language,
  Library,
  Framework,
  Tool
}
relationships { uses, requires, supports }
triples {
  % Foundation.
  System      =requires=> Development;
  Development =uses=> Cocktail;
  % Runtime resources (OS, DBMS, etc.)
  Resource    =supports=> System;
  % Ingredients and their types.
  Ingredient  =pof=> Cocktail;
  Language    =isa=> Ingredient;
  Library     =isa=> Ingredient;
  Framework   =isa=> Ingredient;
  Tool        =isa=> Ingredient;
}

```

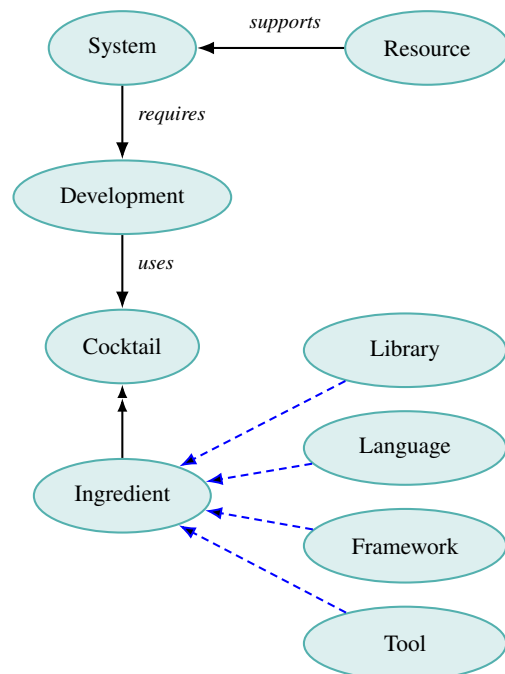


Fig. 3. The inclusion of Ingredients and Resources to the ontology.

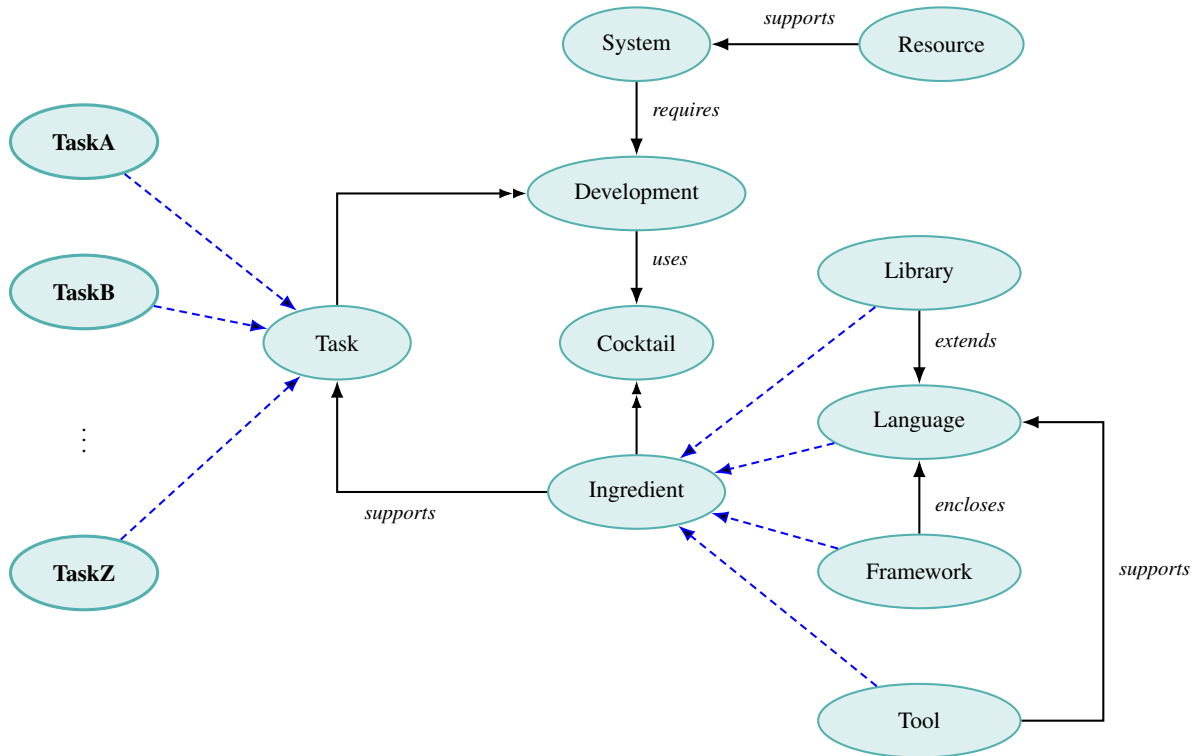


Fig. 4. The conceptual model for the ontology.

heavily on a particular context, there is no optimal solution. A few strategies were considered:

- A general concept is created in order to future-proof the definition, such as *ProgrammingTask*, or equivalent. It does not conceive any particular information about real Tasks that are effectively conducted in the development context. By doing so, the model will not require further updates, while remaining valid for any future Programming Cocktails. The drawback relies on the lack of effective representation, given its generalist nature;
- Given a survey of Programming Cocktails, a set of pre-defined Tasks is established in order to achieve more practical representation than one general concept. These Tasks are fixed and any future use of the ontology will require “fitting” of the actual development context into the set of pre-defined Tasks. This option has the evident risk of rapidly becoming outdated, specially in dynamic domains, such as Web Development. It may also require distortions on the division of Tasks in order to fit the pre-defined concepts, potentially losing its representativeness;
- The conceptual model becomes open and adaptable to specific development contexts. This was the chosen strategy for the ontology. Given that the main intention is to construct knowledge on Programming Cocktails and that each project, team, company, or organization has specific demands and requirements, keeping the ontology adaptable was a better solution to the definition of Tasks.

The final solution for the ontology and how it must deal with Tasks relies on adapting the Task concept (via specialization) to include in the model a logical division of Programming Tasks, fit for the context in case.

As an example, a company with different projects might establish different conceptual models for their ontologies. Suppose that the first project is a simple Web Application, with just a few Ingredients. In this case, the team behind it might simplify the conceptual model and only specialize general Tasks, such as *FrontEndProgramming* and *BackEndProgramming*. In another project, with multiple Ingredients and a much larger problem to solve, the team might find it appropriate to specialize *Task* into a more granular level, such as *LandingPageStructuring*, *ClientInterfaceStyling*, *DatabaseCommunication*, so on, and so forth. This is what the term *open in open conceptual model* means.

This strategy future-proofs the conceptual model by making it adaptable, while providing both flexibility and a solid foundation for Programming Cocktails analysis.

C. Cocktail Identity Cards

The conceptual model of an ontology is crucial to define its structure, how the modelled concepts are related, and what level of detail is expected for the overall organization of knowledge. Nonetheless, the concepts, besides being fundamental, are usually materialized into *individuals*, on their occurrences in the context being modelled.

Listing 4 The conceptual model for the ontology.

```

Ontology Cocktails

concepts {
  System,
  Development,
  Cocktail,
  Resource,
  Ingredient,
  Language,
  Library,
  Framework,
  Tool,
  Task,
  TaskA, TaskB, ..., TaskZ
}

relationships {
  uses,
  requires,
  supports,
  extends,
  encloses
}

triples {
  % Foundation.
  System      =requires=> Development;
  Development =uses=>      Cocktail;
  % Runtime resources (OS, DBMS, etc.)
  Resource    =supports=> System;
  % Ingredients and their types.
  Ingredient  =pof=>      Cocktail;
  Language    =isa=>      Ingredient;
  Library      =isa=>      Ingredient;
  Framework   =isa=>      Ingredient;
  Tool         =isa=>      Ingredient;

  % Language's central role.
  Library      =extends=>  Language;
  Framework    =encloses=> Language;
  Tool         =supports=> Language;
  % General Task concept.
  Task         =pof=>      Development;
  Ingredient    =supports=> Task;
  % Context-specific Tasks. These tasks depend on
  the development context and its structure.
  TaskA        =isa=>      Task;
  TaskB        =isa=>      Task;
  ...
  TaskZ        =isa=>      Task;
}

```

The conceptual model of the ontology was applied to the Cocktails in order to test its validity and aid in structuring the information from the survey. Initially, all individuals and concepts were pictured, which resulted in a convoluted image and many overlapping connections. In order to establish a clearer picture of each Cocktail, the focus shifted to showing the individuals, their relations and, when necessary, some concepts to avoid misidentification of the individuals. The concepts that were kept in the diagram were:

- *Resource* to explicitly show which supporting systems and services each application used;
- *Language*, *Library*, *Framework*, and *Tool*, to categorise

each *Ingredient*;

- *Task* specializations to identify the parts of the application that each *Ingredient* tackles.

In all of these cases, a concept is only added to the diagram if a relation to or from it is also present. As an example, if there are no libraries in the Cocktail, the *Library* concept will not be shown.

The instantiation of an application from the survey is presented in Fig. 5¹⁰. It is a Question & Answer (Q&A) Web Application used for internal communication (and documentation) in the company. It requires three different supporting systems for its execution¹¹: Elasticsearch [44], MongoDB [45], and RabbitMQ [46]. The three basic *Ingredients* for almost any Web Application are present (HTML, CSS, and JavaScript), as are two well known *Frameworks* (Node.js and React.js).

The three *Task* concepts (*WebSiteStyling*, *WebSiteMarkup*, and *WebSiteScripting*) have been determined based on the main areas of the development context. They do not represent specific tasks, as these have not been provided by the company. Their role is to exemplify how *Ingredients* relate to their supported *Tasks*.

As a first proof of concept, the instantiation in Fig. 5 is able to visually represent each development component and how they relate to the application. It does so in a compact form, with enough elements to quickly provide interesting insights into the Cocktail:

- Dependency on external services and systems is directly represented by the number of *Resources* that support the *System*;
- Dependency on *Ingredients* is represented by the number of the equivalent instances;
- Possible redundancies (too many *Ingredients* of the same type supporting the same *Task*) are quickly identified;
- *Tasks* that are too reliant on many *Ingredients*—a possible weakness point—can be directly identified by their number of *support* relations.

The instantiation provides enough information about the application and its development (specially its Cocktail) that it is effectively an *Identity Card* (CIC). It has been successfully applied to the other Cocktails obtained from the survey, providing CICs to all of them. Fig. 6 shows two more CICs for comparison.

The Identity Card shown in Fig. 6a represents a mobile educational game. In this case, the tasks have been chosen in a more granular manner, in order to better represent specific parts of development context. Fig. 6b also represents a mobile application, but not a game. It is a Covid-pass related front-end application. Differently from Fig. 6a, which applied a multi-platform engine (Unity [47]) to create and deploy the game to both mobile application stores (Apple's App Store and Google's Play Store), Fig. 6b shows how one application

¹⁰The names of the applications have been changed to a generic *App#* format for privacy concerns. Nonetheless, they have all been gathered in the survey and represent real software.

¹¹Instantiated from the *Resource* concept.

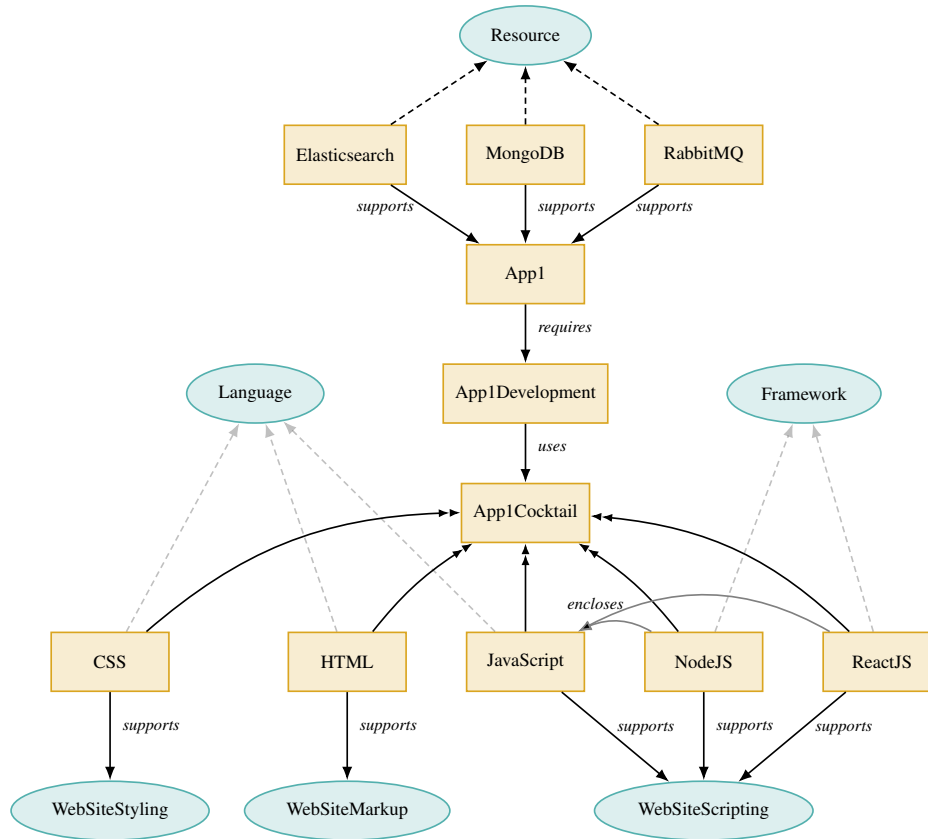


Fig. 5. An example of a Cocktail Identity Card modelled for a Q&A Web Application.

may require more than one development context, since it used mutually exclusive technologies for each platform (Swift for Apple’s ecosystem, and Kotlin for Google’s). In both cases, the CICs quickly present the previously listed properties for their Programming Cocktails.

The Identity Cards are valuable for quick identification of several Cocktail properties, as previously shown, but also form a foundation for the further, deeper analysis. Risks, costs, or any other form of valuation that would be layered on top of their relations could become valid augmentations to the CICs.

D. Structured Knowledge on Cocktails

The instantiations are valuable in their own merit, by organising the relations between the technologies that application development depends on. Nonetheless, its construction, by itself, relayed valuable information about the survey.

The definition of the category columns in the spreadsheet (mentioned in Subsection III-B) is a direct and practical result in this case. The initial version of these columns had several problems, from the lack of domain definitions, to redundancy in values. Since the categorization of the Ingredients will be paramount in future studies, columns *B* to *D* are of great importance. After the definition of the ontology and its application to the several Cocktails that have been gathered, the final version of the category columns was finally obtained.

The first category column (*Type*, column *B*) was a direct implementation of the *Ingredient* specializations (*Language*, *Library*, *Framework* and *Tool*). It is a direct definition of the ingredient’s nature. The next column (*Language*, column *C*) represents what language is used for each other type of ingredient. It was derived directly from the relationships that the different types of ingredients establish to *Language* in the conceptual model (see the bottom-right relations in Fig. 4). Finally, column *D* (*Task*) represents the *Task* specializations, as previously explained. In the case that an ingredient is applied to more than one task, or used with more than one language, its line would be duplicated and changed to reflect these variations.

Another direct result from the construction of the ontology was the possibility to determine which languages are more auto-sufficient (have fewer libraries and frameworks connected to them) or more dependant of complements.

A third point-of-view for knowledge construction based on the ontology relates to project management. A few possibilities include:

- By superimposing the Identity Cards, teams and companies can quickly identify which ingredients they are more dependent on, or have more experience with;
- The definition of the *Task* specializations render an opportunity to identify common threads between projects, in

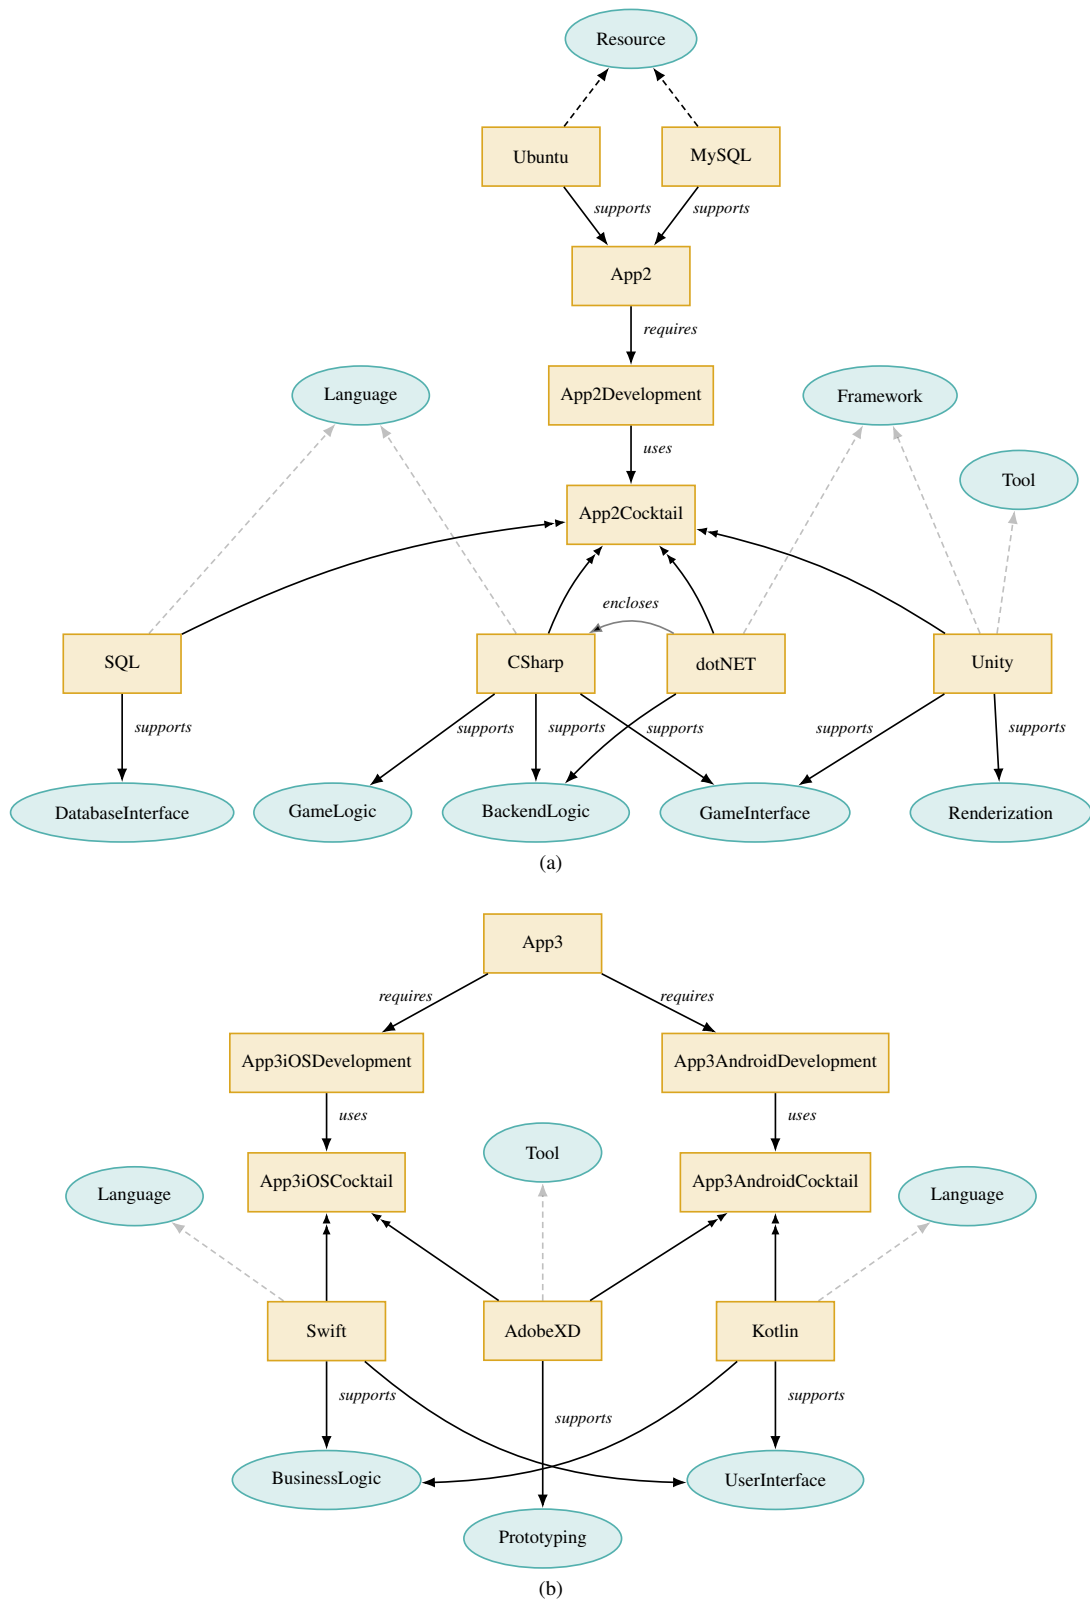


Fig. 6. Identity Cards for two different applications: a mobile educational game (a) and a Covid-pass management front-end (b).

order to standardize or evaluate how teams are structured, personnel is allocated, *etc.*

- The Identity Cards provide quick documentation about a project’s technological evolution. Given that it can also be encoded in OntoDL, it can be easily registered in version control systems, such as Git [48].

As with any kind of structural representation, this ontology may be applied to analyse and support decisions on many facets of project developments, from simple documentation, to critical factors such as risk and dependency.

V. CONCLUSION

This paper presented an ontology based on an open conceptual model of Programming Cocktails. This ontology was initially idealised to aid in defining classifications for ingredients that were surveyed from several companies, and how they related to their development context. This result was successfully accomplished, and a new point-of-view also emerged from the instantiation of the Cocktails: the Identity Cards.

As can be observed from the previous sections, the act of constructing the ontology has already structured knowledge for the surveyed Cocktails, that will provide valuable information for future studies. Nonetheless, the ontology itself and the Identity Cards also presented more interesting opportunities. By creating a foundation on top of which further, more complex analysis could be applied, the Cocktail Identity Cards became a strong result from the ontology. It allows for quick visualization, identification and extraction of knowledge on Programming Cocktails.

Future projects will continue on this development by extending both the methodology for gathering and identifying Cocktails, but also by applying the Identity Cards on consequent studies. For attaining the former, a project for the automatic extraction of Programming Cocktails from public open source projects would be of great use, both for overall analysis, but also as a further proof of concept for the ontology. Also functioning as a direct application of the ontology, the Identity Cards will be augmented with cognitive analysis for a subsequent decision-support system for developers, project managers and teachers.

APPENDIX — SURVEY MESSAGE

The following is a redacted version of the email used to survey Programming Cocktails, as explained in Subsection III-A. Personal details were removed for privacy concerns.

* * *

In the context of our research project, we need to collect information about what we call *Programming Cocktails*, that is, clusters formed by programming languages (marking, formatting, communication...), libraries, frameworks and development environments that are used together to develop complex Applications (software systems).

For example, a very common Cocktail used to develop a current web application is:

- HTML and CSS for formatting;
- JavaScript for client-side programming;

- React as a framework for JS;
- VSCode for development.

As such, we have been contacting prominent companies in the area of Software Development to find out which Cocktails they have effectively used. As you have already shown interest in the academic context, given your participation in our Conferences/Seminars; or even proposing themes for master’s theses, I come to ask: would it be possible for you to reply to this email sending me information about the Programming Cocktails that are used in your company? If this matter is not within your possibilities, we would be extremely grateful if you send us the contact of someone from your company who can help us with this information.

We make it clear that we do not intend to collect any type of confidential information. The answer is informal and should be as simple as shown in the examples below:

Application A, we used:

- Front-end: HTML, CSS, JavaScript, React;
- Back-end: Java, libssh;
- Communication: JSON;
- Database: MySQL;
- Other: Visual Studio, ActiveMQ, ...

Applications B and C, we used:

- Full-stack: Node.js;
- Communication: gRPC;
- Database: MongoDB;
- Other: VSCode, ...

Application D...

We don’t even need the names of the applications, just the list of technologies used in their construction. However, the more details you can provide, such as the area or domain of the application, the more valuable your contribution will be to our research[...]

REFERENCES

- [1] R. R. Fenichel, J. Weizenbaum, and J. C. Yochelson, “A program to teach programming,” *Communications of the ACM*, vol. 13, pp. 141–146, 03 1970. doi: 10.1145/362052.362053. [Online]. Available: <https://dl.acm.org/doi/10.1145/362052.362053>
- [2] M. J. V. Pereira and P. R. Henriques, “Visualization/animation of programs in alma: Obtaining different results,” in *Proceedings of the IEEE Symposium on Human Centric Computing Languages and Environments*, 2003. doi: 10.1109/HCC.2003.1260242 pp. 260–262. [Online]. Available: <https://ieeexplore.ieee.org/document/1260242>
- [3] T. C. Freitas, A. Costa Neto, M. J. V. Pereira, and P. R. Henriques, “Nlp/ai based techniques for programming exercises generation,” R. A. P. d. Queirós and M. P. T. Pinto, Eds., vol. 104, Open Access Series in Informatics (OASICS). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi: 10.4230/OASICS.SLATE.2022.14 pp. 1–15. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2022/16760>
- [4] S. A. Teixeira, “Automatic grading of programming exercises,” Master’s thesis, Minho University, Braga, Portugal, 2023, to be published.
- [5] P. Vasconcelos, “Haskellite: A step-by-step interpreter for teaching functional programming,” R. A. P. d. Queirós and M. P. T. Pinto, Eds., vol. 104, Open Access Series in Informatics (OASICS). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi: 10.4230/OASICS.SLATE.2022.14 pp. 1–15. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2022/16760>
- [6] S. A. Robertson and M. P. Lee, “The application of second natural language acquisition pedagogy to the teaching of programming languages: a research agenda,” *ACM SIGCSE Bulletin*, vol. 27, no. 4, p. 9–12, 12 1995. doi: 10.1145/216511.216517. [Online]. Available: <https://dl.acm.org/doi/10.1145/216511.216517>
- [7] M. V. P. Almeida, L. M. Alves, M. J. V. Pereira, and G. A. R. Barbosa, “Easycoding: Methodology to support programming learning,” R. Queirós, F. Portela, M. Pinto, and A. Simões, Eds., vol. 81, Open Access Series in Informatics (OASICS). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 06 2020. doi: 10.4230/OASICS.ICPEC.2020.1. ISBN 978-3-95977-153-5. ISSN 2190-6807 pp. 1–8. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2020/12288>

- [8] J. L. Plass, B. D. Homer, and C. K. Kinzer, "Foundations of game-based learning," *Educational Psychologist*, vol. 50, no. 4, pp. 258–283, 2015. doi: 10.1080/00461520.2015.1122533
- [9] A. Gomes and A. J. Mendes, "Learning to program: Difficulties and solutions," Proceedings of the 2007 International Conference on Engineering and Education (ICEE). International Network on Engineering Education and Research, 2007, pp. 283–287. [Online]. Available: <http://icee2007.dei.uc.pt/proceedings/papers/411.pdf>
- [10] B. C. Wilson and S. Shrock, "Contributing to success in an introductory computer science course: a study of twelve factors," Proceedings of the 32nd SIGCSE Technical Symposium on Computer Science Education. Association for Computing Machinery, 2001. doi: 10.1145/364447.364581. ISBN 1581133294 pp. 184–188. [Online]. Available: <https://dl.acm.org/doi/10.1145/364447.364581>
- [11] P. C. Tavares, E. M. F. Gomes, and P. R. Henriques, "O impacto da animação e da avaliação automática na motivação para o ensino da programação," Ph.D. dissertation, 2017.
- [12] A. Costa Neto, C. Araújo, M. J. V. Pereira, and P. R. Henriques, "Programmers' affinity to languages," P. R. Henriques, F. Portela, R. Queirós, and A. Simões, Eds., vol. 91, Open Access Series in Informatics (OASICs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi: 10.4230/OASICs.ICPEC.2021.3. ISBN 978-3-95977-194-8. ISSN 2190-6807 pp. 1–7. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2022/16760>
- [13] J. Alves, A. Costa Neto, M. J. V. Pereira, and P. R. Henriques, "Characterization and identification of programming languages," A. Simões, M. M. Berón, and F. Portela, Eds., vol. 104, Open Access Series in Informatics (OASICs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi: 10.4230/OASICs.SLATE.2022.14 pp. 1–15. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2022/16760>
- [14] C. Casalnuovo, E. T. Barr, S. K. Dash, P. Devanbu, and E. Morgan, "A theory of dual channel constraints," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: New Ideas and Emerging Results*, ser. ICSE-NIER '20. New York, NY, USA: Association for Computing Machinery, 2020. doi: 10.1145/3377816.3381720. ISBN 9781450371261 p. 25–28. [Online]. Available: <https://doi.org/10.1145/3377816.3381720>
- [15] J. Figueiredo and F. J. García-Peñalvo, "Building skills in introductory programming," F. J. García-Peñalvo, Ed., Proceedings of the Sixth International Conference on Technological Ecosystems for Enhancing Multiculturality. New York: ACM, 10 2018. doi: 10.1145/3284179. ISBN 9781450365185 p. 46–50. [Online]. Available: <https://dl.acm.org/doi/10.1145/3284179.3284190>
- [16] M. Fourment and M. R. Gillings, "A comparison of common programming languages used in bioinformatics," *BMC Bioinformatics*, vol. 82, no. 9, 02 2008. doi: 10.1186/1471-2105-9-82. [Online]. Available: <https://bmcbioinformatics.biomedcentral.com/articles/10.1186/1471-2105-9-82>
- [17] A. H. Odeh, "Analytical and comparison study of main web programming languages: Asp and php," *TEM Journal*, vol. 8, pp. 1517–1522, 11 2019. doi: 10.18421/TEM84-58. [Online]. Available: http://www.temjournal.com/content/84/TEMJournalNovember2019_1517_1522.pdf
- [18] N. Walia and A. Kalia, "Programming languages for data mining: a review," *International Journal of Computer Trends and Technology*, vol. 68, pp. 38–41, 2020. doi: 10.14445/22312803/IJCTT-V68I1P109. [Online]. Available: <https://ijcttjournal.org/archives/ijctt-v68i1p109>
- [19] T. R. Gruber, "A translation approach to portable ontology specifications," *Knowledge Acquisition*, vol. 5, no. 2, pp. 199–220, 1993. doi: <https://doi.org/10.1006/knac.1993.1008>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1042814383710083>
- [20] R. Studer, V. R. Benjamins, and D. Fensel, "Knowledge engineering: Principles and methods," *Data & Knowledge Engineering*, vol. 25, no. 1, pp. 161–197, 1998. doi: [https://doi.org/10.1016/S0169-023X\(97\)00056-6](https://doi.org/10.1016/S0169-023X(97)00056-6). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169023X97000566>
- [21] Apache Software Foundation, "Apache http server project." [Online]. Available: <https://httpd.apache.org>
- [22] —, "Apache activemq." [Online]. Available: <https://activemq.apache.org>
- [23] Oracle, "MySQL." [Online]. Available: <https://www.mysql.com>
- [24] Dormando, "Memcached." [Online]. Available: <https://www.memcached.org>
- [25] B. W. Kernighan and D. M. Ritchie, *C Programming Language*, 2nd ed. Pearson, 03 1988.
- [26] Python Foundation, "Welcome to python.org," 11 2019. [Online]. Available: <https://www.python.org>
- [27] Mozilla Foundation, "Html: Hypertext markup language." [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/HTML>
- [28] —, "Css: Cascading style sheets." [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/CSS>
- [29] D. D. Chamberlin, "Early history of sql," *IEEE Annals of the History of Computing*, vol. 34, pp. 78–82, 11 2012. doi: 10.1109/MAHC.2012.61. [Online]. Available: <https://ieeexplore.ieee.org/document/6359709>
- [30] LibSSH, "Libssh." [Online]. Available: <https://www.libssh.org>
- [31] R. Santamaria, "raylib." [Online]. Available: <https://www.raylib.com>
- [32] Apple, "Swiftui," Apple Developer. [Online]. Available: <https://developer.apple.com/xcode/swiftui/>
- [33] Meta Platforms, "React native." [Online]. Available: <https://reactnative.dev>
- [34] The JUnit Team, "JUnit." [Online]. Available: <https://junit.org>
- [35] C. Araújo, P. R. Henriques, and J. J. Cerqueira, "Ontocne, characterizing learning resources for training computational thinking," in *2023 International Symposium on Computers in Education (SIIE)*, 2023. doi: 10.1109/SIIE59826.2023.10423710 pp. 1–6.
- [36] S. Teixeira, R. V. Boas, F. Oliveira, C. Araújo, and P. R. Henriques, "Ontojogo: An ontology for game classification," in *2020 IEEE 8th International Conference on Serious Games and Applications for Health (SeGAH)*. Vancouver, BC, Canada: IEEE Xplore, 2020. doi: 10.1109/SeGAH49190.2020.9201876. ISBN 978-1-7281-9042-6. ISSN 2573-3060 pp. 1–8.
- [37] C. Araújo, L. Lima, and P. R. Henriques, "An Ontology based approach to teach Computational Thinking," in *21st International Symposium on Computers in Education (SIIE)*, C. G. Marques, I. Pereira, and D. Pérez, Eds. IEEE Xplore, Nov 2019. doi: <https://doi.org/10.1109/SIIE48397.2019.8970131>. ISBN 978-1-7281-3182-5 pp. 1–6.
- [38] D. R. Barbosa, "CnE-Ar: Teaching of Computational Thinking to Adults in Reconversion," Master's thesis, Minho University, Braga, Portugal, April 2021, MSc dissertation.
- [39] M. de La Saleta Teixeira, "Adequa, a platform for choosing Games suitable to Students' Profile," Master's thesis, Minho University, Braga, Portugal, March 2021, MSc dissertation.
- [40] C. Araújo, P. R. Henriques, and J. J. Cerqueira, "Creating Learning Resources based on Programming concepts," in *Local Proceedings of the 15th International Conference on Informatics in Schools – ISSEP 2022*, A. Bollin and G. Futschek, Eds. Klagenfurt; Wien, Austria: The Austrian Library Association, open-access net-library, Sep 2022. doi: 10.48415/2022/issep.2022 pp. 35–46.
- [41] L. Martins, C. Araújo, and P. R. Henriques, "Digital Collection Creator, Visualizer and Explorer," in *8th Symposium on Languages, Applications and Technologies (SLATE 2019)*, ser. OpenAccess Series in Informatics (OASICs), R. Rodrigues, J. Janoušek, L. Ferreira, L. Coheur, F. Batista, and H. G. Oliveira, Eds., vol. 74. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. ISBN 978-3-95977-114-6 p. 15:1–15:8. [Online]. Available: <https://www.dagstuhl.de/dagpub/978-3-95977-114-6>
- [42] A. M. C. Dias, "ONTODL+, An ontology description language and its compiler," Master's thesis, Minho University, Braga, Portugal, Sep 2021, MSc dissertation.
- [43] W3C, "Web ontology language (owl)." [Online]. Available: <https://www.w3.org/OWL/>
- [44] Elastic, "Elasticsearch." [Online]. Available: <https://www.elastic.co/elasticsearch>
- [45] MongoDB, "Mongodb." [Online]. Available: <https://www.mongodb.com>
- [46] Broadcom, "Rabbitmq." [Online]. Available: <https://www.rabbitmq.com>
- [47] Unity Technologies, "Unity real-time development platform." [Online]. Available: <https://unity.com>
- [48] Software Freedom Conservancy, "Git." [Online]. Available: <https://git-scm.com>

SrpCNNeL: Serbian Model for Named Entity Linking

Milica Ikonić Nešić^{*✉}, Saša Petalinkar^{†✉}, Ranka Stanković^{‡✉}, Miloš Utvić^{*✉} and Olivera Kitanović^{‡✉}

^{*}University of Belgrade, Faculty of Philology, Belgrade, Serbia

Email: milica.ikonic.nesic@fil.bg.ac.rs, milos.utvic@fil.bg.ac.rs

[†]University of Belgrade, Belgrade, Serbia

Email: sasa5linkAr@gmail.com

[‡]University of Belgrade, Faculty of Mining and Geology, Belgrade, Serbia

Email: ranka.stankovic@rgf.bg.ac.rs, olivera.kitanovic@rgf.bg.ac.rs

Abstract—This paper presents the development of a Named Entity Linking (NEL) model to the Wikidata knowledge base for the Serbian language, named SrpCNNeL. The model was trained to recognize and link seven different named entity types (persons, locations, organizations, professions, events, demonyms, and works of art) on a dataset containing sentences from novels, legal documents, as well as sentences generated from the Wikidata knowledge base and the Leximirka lexical database. The resulting model demonstrated robust performance, achieving an F1 score of 0.8 on the test set. Considering that the dataset contains the highest number of locations linked to the knowledge base, an evaluation was conducted on an independent dataset and compared to the baseline Spacy Entity Linker for locations only.

I. INTRODUCTION

NAMED Entity Linking (NEL) is one of the important tasks in Natural Language Processing (NLP) that focuses on identifying and disambiguating entities mentioned in the text by linking them to a corresponding knowledge base [1]. This process is essential for structuring unstructured data, which is crucial for various NLP applications such as information retrieval, sentiment analysis, and knowledge base population. NEL is particularly significant for low-resource, morphologically rich languages like Serbian due to the unique challenges and benefits it presents [2], [3].

During the last decade, various works have addressed named entity linking, recognition, and disambiguation. The task of recognizing mentions of entities in text and disambiguating them to the corresponding entities in a knowledge base (KB) is called Entity Linking (EL)[4]. EL systems have demonstrated remarkable performance on standard benchmarks, a feature largely attributable to the advent of contemporary language models[5]. It has found innumerable applications in a wide range of downstream tasks, such as Question Answering [6], Information Extraction [7], Historic Newspaper optical character recognition (OCR) [8], Esports News [9], EL for Tweets [10], and Biomedical areas [11], [12]. The importance of NEL extends beyond entity recognition; it plays a crucial role in organizing and extracting meaningful information from large text corpora. Effective NEL can significantly improve the accuracy of downstream NLP tasks by providing a structured understanding of entities and their relationships within the

text. This is especially important for processing large volumes of unstructured data, which is common in many real-world applications. For instance, NEL can improve information retrieval by linking query terms to relevant entities, enhance sentiment analysis by accurately identifying sentiment targets, and support KB population by automatically updating entity information.

New research focuses on a survey of the scientific literature on NEL, including named entity recognition and disambiguation, covering 200 works by focusing on 43 papers (5 surveys and 38 research works). The authors also described and classified 56 resources, including 25 tools and 31 corpora for English and other languages [13]. Balog [1] defined the problems of EL, NER, and NED (Named Entity Disambiguation) by presenting the general process of Entity Linking (EL), which consists solely of NER and NED. He asserts that named entities must be extracted before they can be disambiguated.

ScispaCy [14] integrates spaCy for biomedical text processing, offering efficient and accurate entity linking to knowledge bases like the Unified Medical Language System (UMLS)[15] and Gene Ontology[16]. These models demonstrate significant improvements in linking biomedical entities, highlighting spaCy's versatility in handling domain-specific NLP tasks. Radboud Entity Linker (REL) [17] is an open-source toolkit for entity linking, which builds on neural components from natural language processing research and is provided as a Python package and a web API.

The linking of named entities has sharply increased over the years for multilingual models [18], [19], [20]. Bi-encoder Entity Linking Architecture (BELA) presents the first fully end-to-end multilingual entity linking model that efficiently detects and links entities in texts in any of 97 languages [21]. Additionally, there has been a growing body of research focused on enhancing neural models by leveraging relational knowledge from semantic networks using Wikidata [22], [23], [24]. Linking entities to the Wikidata knowledge base is currently a highly relevant topic, and one of the projects from 2021 is the Spacy Entity Linker (spacy-entity-linker 1.0.3) pipeline for spaCy that performs linked entity extraction with Wikidata, which can be used as a multilingual entity model for linking with the Wikidata knowledge base. The University

Library of Mannheim (Universitätsbibliothek Mannheim, abbreviated UB Mannheim) developed spaCyOpenTapioca for the task of linking named entities to concepts (items) in Wikidata in spaCy using OpenTapioca [25]. This system achieved an F1 score of 0.09 on an Italian-Serbian corpus of 10,000 aligned segments (sentences) taken from different novels, named the It-Sr-NER corpus [26], [27]. Early research leveraged word2vec and convolutional neural networks (CNN) to capture the correlation between mentioned context and entity information [28], [29] and link entities for languages such as Chinese [30] and Italian [31]. An end-to-end entity linking system for the Greek language was developed in order to extend the Radboud Entity Linker (REL) toolkit to support modern Greek. The authors investigate three different mention detection approaches using spaCy, Flair, and BERT [32].

However, one common issue with current EL approaches is that they require massive amounts of training data. Consequently, training models for named entity recognition and linking with knowledge bases for low-resource languages is a highly challenging endeavor. In the case of the Serbian language, the problem encountered when applying multilingual models for linking recognized entities to the Wikidata knowledge base is the inability to recognize inflected forms of entities. Resolving this problem was one of the main motivations for this research.

Serbian, with its complex morphological structures and limited NLP resources, poses significant hurdles for accurate named entity recognition and linking. The intricate morphology, including rich inflectional patterns, necessitates advanced models capable of handling various word forms and syntactic nuances. Implementing NEL for Serbian using spaCy provides a robust framework to address these challenges by leveraging advanced machine-learning techniques and pre-trained language models tailored for low-resource settings. Considering the challenges presented by the Serbian language, this work represents, to the best of our knowledge, one of the first attempts to train a model for named entity recognition and linking to the corresponding items in Wikidata, with a primary focus on locations using the previously trained SrpCNER2 model for the NER task. This model was trained on a dataset containing Serbian novels published between 1840 and 1920 [33], publicly available newspaper articles, and sentences generated for the NER task from the Wikidata [34] knowledge base and Leximirka lexical database [35], achieving an F1 score of approximately 0.71 on the test dataset.

In conclusion, implementing named entity linking for Serbian using spaCy addresses the specific linguistic challenges posed by the language's morphology and enhances the overall quality and usability of NLP applications in low-resource settings. This research underscores the importance of developing tailored NLP solutions that cater to the unique needs of morphologically rich languages, ultimately contributing to more inclusive and comprehensive language technologies. The paper is organized into several sections as follows. Section II briefly presents the process of preparing the dataset for the training model. Data conversion and the development of a

knowledge base, which enabled the training of a CNN-based NEL model for Serbian, named SrpCNeL, are presented in Section III. Evaluation of the model in two different settings: the first discusses the model's performance on the test subset of the prepared dataset, and the second carries out a detailed evaluation on novel and newspaper articles that were not present in the training dataset, can be found in Section IV. Finally, conclusions and plans for future work are presented in Section V.

II. DATA PREPARATION

One of the key issues addressed in this paper was overcoming the problem of linking inflected forms of entities mentioned in the text with the Wikidata KB. Since items in Wikidata have labels in the nominative form, such as *Beograd* (Q3711) (Belgrade), it was necessary to create a synthetic dataset with inflected forms to ensure that all inflected forms of the mentioned entity in the text are linked with the same item in the KB. For instance, *Beograda*, *Beogradu*, or *Beogradom* are linked to the item *Beograd* (Q3711).

Therefore, for the purposes of this research, the training dataset was created and evaluated in two stages. In the first stage, the synthetic training dataset was generated using data from the Wikidata knowledge base and the Leximirka lexical database [35].

Given that the lexical database Leximirka has semantic markers such as *NPropN* for named entities, as well as **Dr** for country, **Gr** for city, **Hum** for person, **Oro** for mountain, **Hyd** for river, and **Org** for organizations, it was possible to access named entities that belong to any of the mentioned categories. On the other hand, using the Wikidata Query Service and constructing a SPARQL query [36], it was possible to create a list of all items belonging to the categories of country, city, mountain, river, or organization which have labels in the Serbian language. This was achieved by using specific properties in Wikidata for each category, as shown in Table I.

The SPARQL query for extracting instances of countries with appropriate QIDs from Wikidata can be retrieved using the following query:

```
# Countries with labels in Serbian
SELECT DISTINCT ?country ?cLabel
WHERE {
  ?country rdfs:label ?cLabel;
  #instance of country
  wdt:P31 wd:Q6256.
filter langMatches(lang(?cLabel), "sr")
}}
```

TABLE I
PROPERTIES FOR SPECIFIC CATEGORY OF NAMED ENTITIES

semantic marker	category	instance of (P31)
Dr	country	wd:Q6256
Gr	sity	wd:Q515
Oro	mountain	wd:Q46831
Hyd	river	wd:Q4022

TABLE II
EXAMPLE OF SQL FUNCTIONS AND GENERATED SYNTHETIC SENTENCES

mark	num.	SQL function	synthetic sentences
Gr	4	select * from dbo.fnGenerisiReceniceNER(N'<s>Posetio sam <LOC>',N'</LOC> prošlog leta. </s>';'Gr';4',100)	('Posetio sam Beograd prošlog leta. ' 'links': (12,19): 'Q3711':1.0,'entities':{(12,19,'LOC')})
Dr	2	select * from dbo.fnGenerisiReceniceNER(N'<s>U srcu <LOC>',N'</LOC> leži bogatstvo narodnih običaja. </s>';'Dr';2',100)	('U srcu Srbije leži bogatstvo narodnih običaja. ' 'links': (7,13): 'Q403':1.0,'entities':{(7,13,'LOC')})
Hyd	2	select * from dbo.fnGenerisiReceniceNER(N'<s>Voda <LOC>',N'</LOC> je bistra. </s>';'Hyd';2',60)	('Voda Dunava je bistra. ' 'links': (5,11): 'Q1653':1.0,'entities':{(5,11,'LOC')})
Oro	7	select * from dbo.fnGenerisiReceniceNER(N'<s>Našli smo izvor vode na <LOC>',N'</LOC>. </s>';'Oro';5',60)	('Našli smo izvor vode na Durmitoru . ' 'links': (24,33): 'Q212836':1.0,'entities':{(24,33,'LOC')})

After extracting QIDs from Wikidata, it was possible to input them into Leximirka and link them to the corresponding lexical units. Since Leximirka contains inflected forms of all lemmas, it was possible to derive the correct grammatical forms for all grammatical cases of named entities that appear in it and that are linked to Wikidata. A function was developed in MS SQL Server to generate synthetic sentences for a given sentence template. In Table II, examples of SQL queries and appropriately generated synthetic sentences are presented. The column "num" represents the ordinal number of the case in the Serbian language.

In the end, this dataset contains 16,869 sentences, including only those with locations (rivers, cities, countries, lakes, and mountains) and organizations.

After creating a synthetic dataset, the dataset was expanded by 22,730 sentences from various novels written or translated into Serbian, as well as legal documents. The novels whose sentences belong to this dataset include Jules Verne's "Around the World in Eighty Days" [37], Orwell's "1984" [38], and novels from the It-Sr-NER corpus [26]: Umberto Eco's "The Name of the Rose", Carlo Collodi's "The Adventures of Pinocchio", Elena Ferrante's "Those Who Leave and Those Who Stay", and Luigi Pirandello's "One, None and a Hundred Thousand". The corpus also includes five novels by Serbian writers: Ivo Andrić's "Anikina vremena" ("Legends of Anika") and "Na drini ćuprija" ("The Bridge on the Drina"), Borisav Stanković's "Nečista krv" ("Impure Blood"), as well as legal documents from *Intera* available on the Bilibish digital library [39]. The training dataset was prepared through several steps. Initially, the dataset was annotated with named entities using the Jerteh-355-tesla [40] and the Jerteh-355 [41] fine-tuned language model for NER. Jerteh-355 is based on the RoBERTa-large architecture [42]. The model is trained to recognize seven categories: demonyms (DEMO), professions and titles (ROLE), works of art (WORK), person names (PERS), locations (LOC), events (EVENT), and organizations (ORG). After automatic annotation, the INCEpTION tool [43] was used for the manual correction and linking of named entities with the Wikidata knowledge base. An example of annotation and linking with Wikidata using the INCEpTION tool is presented in Fig. 1. The named entities were linked in an additional layer of annotation where a Wikidata identifier was assigned to each entity instance.

After the dataset was prepared, it contained 35,955 sen-

tences in total. Among the annotated entities, as well as those linked to Wikidata, the majority were recognized as locations (LOC), of which only 322 were not linked to Wikidata (Fig. 2), primarily those locations for which corresponding items do not exist in Wikidata.

The distribution of all named entities by type, showing the proportion of those linked to Wikidata versus to those that are not, as well the number of all named entities of appropriate class, is illustrated in the Fig. 3 The Fig. 4 shows the percentage of unique entities relative to the total number of entities recorded in the dataset. Observing the further distribution of linked entities by the number of appearances in the dataset, Serbia (Q403) is the most frequently linked entity to Wikidata. Table III presents the ten entities that are most frequently linked to Wikidata. Although such a distribution is not optimal for training entity linking in a morphologically rich language, it is not uncommon for a corpus dominated by natural text. Fig. 5 illustrates the relationship between the number of entity occurrences and the frequency of entity repetitions. It can be observed that nearly 3,000 entities appear only once in the dataset, while only one entity (Serbia (Q403)) appears more than 700 times.

III. SRPCNNEL MODEL FOR SERBIAN

This section outlines the methodology employed to train the Entity Linker for the Serbian language using the spaCy framework. The process encompassed data preparation, the de-

TABLE III
TEN MOST COMMON LINKED ENTITIES IN DATASET

QID	Name	Count
Q403	Srbija	785
Q236	Crna Gora	403
Q11428966	Vinston Smit (1984)	327
Q37226	učitelj	277
Q838261	Savezna Republika Jugoslavija	248
Q2587533	Fileas Fog	243
Q127885	Srbi	239
Q327055	radnik	222
Q170287	Subotica	188
Q37024	Srbija i Crna Gora	183

velopment of a knowledge base using Wikidata, the training of the entity linking model, and the evaluation of its performance.

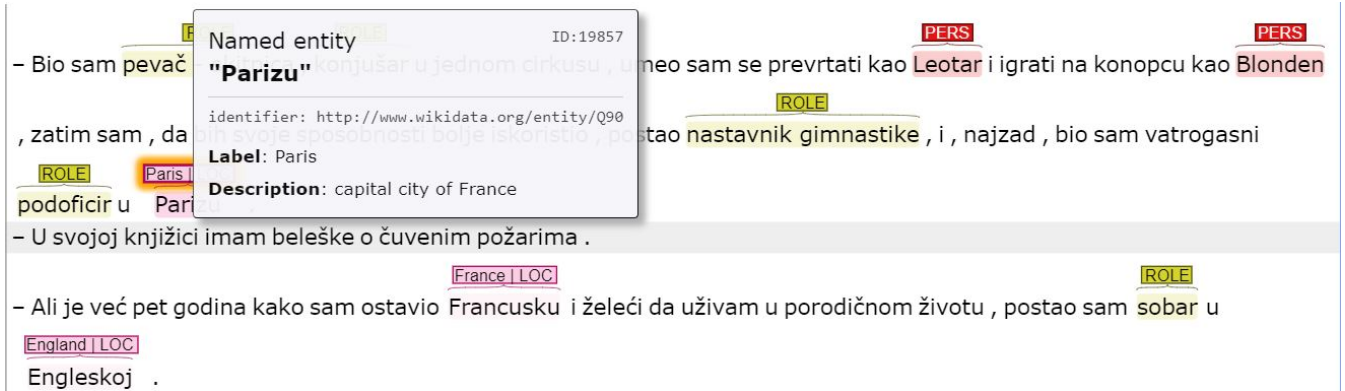


Fig. 1. An example of annotation in INCEpTION

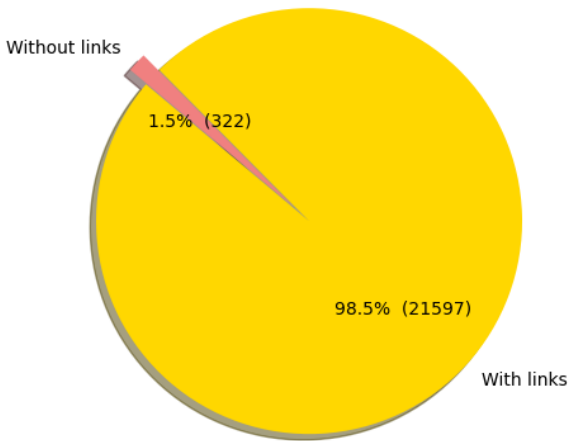


Fig. 2. Percentage of entities that are linked to the KB and those that are not




Fig. 4. Percent of unique linked entities in corpus




Fig. 3. Distribution of named entity types with and without linking to Wikidata




Fig. 5. Distribution of unique entities by the number of appearance in the dataset

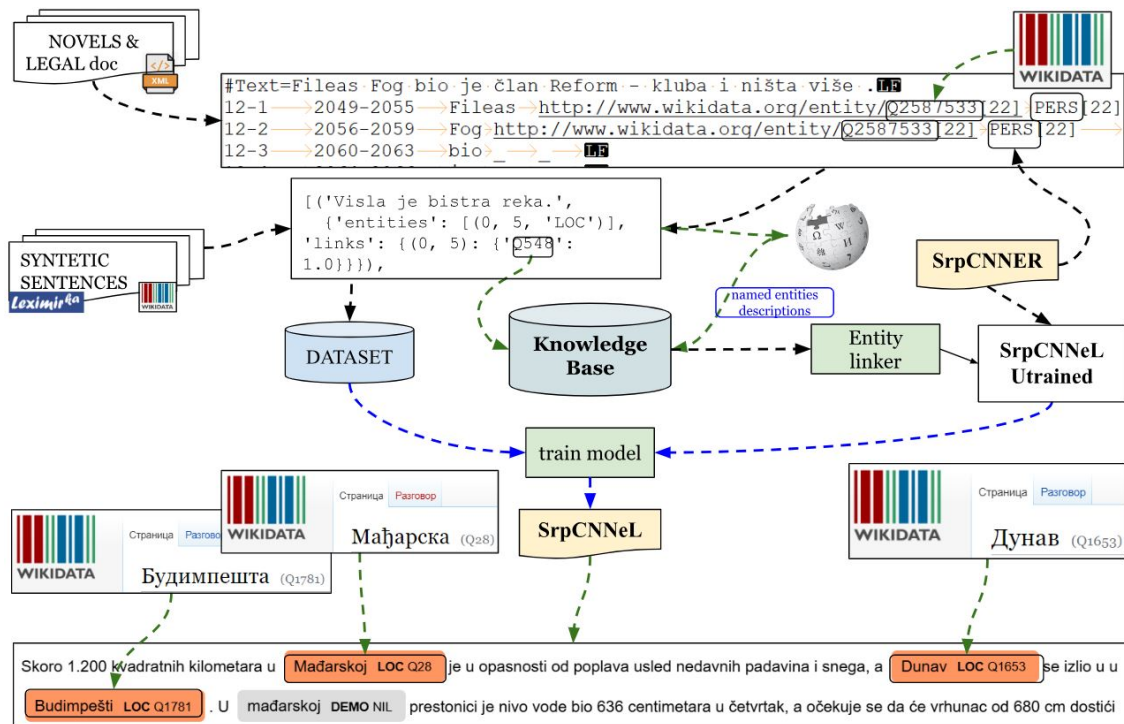


Fig. 6. The workflow of the whole process of NEL

The workflow of the previous process is presented in Fig. 6. evaluation of its performance.

A. Data Conversion

For training the NEL model, the training dataset contained 35,955 sentences in total. The research involved converting data from different formats into one usable by spaCy, with serialization into the .pkl (pickle) format. The pickle format is efficient for storage and quick loading during model training and inference. The code used for this research included functions to convert TSV and custom text formats into a spaCy-compatible format, ensuring accurate entity linking and recognition.

B. Knowledge Base Development

The knowledge base was developed from the extracted entities and appropriate QIDs in the dataset and enriched with descriptions obtained from Serbian Wikipedia pages. This involved extracting entities and their corresponding descriptions to form a comprehensive knowledge base. This step was particularly important for morphologically rich languages like Serbian, as it allowed for the capturing of all inflected forms of words that appear in the dataset.

Instead of relying on lemmatization, inflected forms of words were incorporated into the knowledge base. This decision stems from the complexity of lemmatizing multi-word named entities in Serbian, where a token-by-token lemmatization approach is inadequate and often incorrect.

Unique identifiers from Wikidata were extracted for each entity, along with all texts that correspond to each entity's aliases. Probabilities were not extracted from the dataset; instead, each alias was assigned a uniform starting probability. Descriptions from Wikipedia were connected to each unique identifier from Wikidata and vectorized using the base spaCy pipeline. These descriptions of entities, as well as all possible aliases, were then inserted into the knowledge base. Ensuring that inflected forms are included in the knowledge base facilitates their recognition, making it possible to accurately link entities in the text to their corresponding entries in the Wikidata knowledge base. This comprehensive approach ensures that the entities recognized in the text are accurately linked, enhancing the model's performance in handling the complex linguistic variations inherent in the Serbian language.

C. Training the Entity Linking Model

The SrpCNNeL model extends the pre-trained NER model for Serbian, SrpCNNeL2, by integrating a spaCy entity linking layer, where the base model employed for entity linking is `spacy.EntityLinker.v2`. The SrpCNNeL2 model is trained using the spaCy Python module, version 3.2, employing the same model architecture as SrpCNNeL [44], using 148,819 sentences dataset, sourced from three distinct components: `srpELTeC-gold-extended`, newspaper articles and generated sentences. Importantly, the dataset includes sentences that do not contain any named entities. The breakdown is as follows:

- **srpELTeC-gold-extended**: This component contributes 54,423 sentences, encompassing 986,567 tokens and

35,772 named entities. This corpus is an extension of the srpELTeC-gold corpus [45], which contains sentences from old Serbian novels within the SrpELTeC corpus.

- **Newspaper articles:** This segment consists of 9,498 sentences, amounting to 235,953 tokens and 28,496 named entities.
- **Generated sentences:** This category comprises the largest portion, with 84,898 sentences, totaling 670,722 tokens and 85,642 named entities. Sentences are generated by the Wikidata knowledge base and sentences generated on the basis of Leximirka lexical database [46], in the form of CONLLu files which contain information about NER, POS-tag, and lemma.

SrpCNER2 is trained to recognize seven categories of entities: persons (PERS), professions (ROLE), demonyms (DEMO), organizations (ORG), locations (LOC), artworks (WORK), and events (EVENT). The SrpCNeL model was trained using the previously converted data in pkl format and the created knowledge base. Sentences were randomly shuffled and split into training and test sets with a ratio of 8:2, i.e., 31,679 sentences in training and 3,720 sentences in the test set. SpaCy's pipeline was employed to integrate the NER component with the *entity linker*. To feed training data into the entity linker, the pkl format presents a list of structured tuples. The first part is the raw text, and the second part is a dictionary of annotations. The dictionary defines the named entities we want to link ("entities"), as well as the actual gold-standard links ("links"). An example of such a tuple is the following:

```
('Visla je bistra reka.', 'links': (0, 5): 'Q548': 1.0, 'entities': [(0, 5, 'LOC')])
```

This integration was essential to ensure that entities identified in the text were correctly linked to their corresponding entries in the knowledge base. The dropout rate was set to 0.2, and the optimizer was inherited from SrpCNER2. The training data, serialized into pickle files, was loaded, and the model was trained iteratively (with 100 iterations) using examples from the training dataset.

IV. EVALUATION

Model evaluation performed on the previously discussed test set demonstrated the following performance metrics: a precision of 0.79, a recall of 0.83, and an F1 score of 0.80, indicating a robust model performance.

In addition to this, and in order to better understand how the model behaves at the sentence level, we introduce a new metric called "*accuracy by sentence*". This measure provides a more holistic view of the model's performance by considering the correct prediction of entire sentences rather than individual entities. A sentence is considered correct if all predicted entities in it correspond to those annotated in the dataset, pointing to the same text and being linked identically. The type of entity was not considered, as the spaCy entity linker does not account for entity types. Accuracy by sentence is calculated according to Equation 1.

$$accuracy_{sentence} = \frac{correct_{sentence}}{total_{sentences}} \quad (1)$$

Introducing this metric is motivated by the need to evaluate the model's performance in practical, real-world scenarios, where accurate entity linking across entire sentences is crucial for applications such as information extraction, knowledge base population, and automated content analysis. By focusing on sentence-level accuracy, we can better assess the model's ability to understand and process context, ensuring that linked entities are not only identified correctly but also coherent within their respective sentences. The result on the test set is an accuracy by sentence of 0.67, as shown in Fig. 7. This highlights the model's capability to correctly link entities in a significant portion of the sentences, though there remains room for improvement to achieve higher accuracy in more complex or nuanced cases.

In this research, the main focus was on the recognition of linked locations and organizations, as the occurrence of names and roles depends on the text itself. For example, literary texts more frequently feature *characters* from novels or *historical persons*, while newspapers predominantly mention *politicians*, *athletes*, *actors*, etc. Even the roles differ, with newspapers mentioning *politicians*, *collaborators*, *directors*, while in the novels within this corpus, roles such as *duke*, *king*, *servant*, etc., are more prevalent. However, to highlight both the advantages and disadvantages of such a system, in this study, all named entities present in the knowledge base were linked within the corpus.

The performance achieved by the model indicates that a larger dataset is necessary for training the entity linking model in this case. The scatter plot presented in Fig. 8 suggests that the model has generally good classification performance, as




Fig. 7. Accuracy by sentence




Fig. 8. Scatter plot by Wikidata links

indicated by the concentration of points along the diagonal line. The vertical line of points at a specific true value signifies that there is an entity with a high frequency and a high degree of prediction error. This entity likely has a large number of members, making it challenging for the model to predict accurately.

On closer examination, this entity turned out to be one marked NIL, which is a catch-all group of all non-linked entities.

The scatter of points away from the diagonal line indicates prediction errors, suggesting that while the model is effective overall, there are certain cases where it does not perform as well.

A. Separate Evaluation Set

This section presents the evaluation on an independent dataset only for locations. For this purpose, 299 sentences from the novel "The Good Soldier Švejk" by Jaroslav Hašek and 510 sentences from the newspaper "Politika" have been prepared. The evaluation results will be demonstrated by comparing the newly trained model SrpCnNeL with the custom Spacy Entity Linker. We took the strictest approach and differentiated between the following three situations:

- [TP] an entity is recognized exactly as it should, comparing to the gold standard (the text and the QID match – true positives);
- [FP] an entity is recognized, but not with the correct QID;
- [FN] an entity present in the gold standard was not recognized.

The results for the SrpCnNeL model are displayed in the upper part of Table IV.

In the case of the newspaper article, a larger number of entities not linked to a QID but present in the gold standard (FN) are attributed to the entity named Ukraine and all its forms in different grammatical cases, except for the form "Ukrajinu," which is recognized. Upon deeper analysis, it was determined that, at the time of extraction of QIDs, Ukraine did not have the property "instance of state" on Wikidata but instead had other properties such as sovereign state (Q3624078), social state (Q619610), and territory (Q4835091). Consequently, it was not included in the synthetic dataset.

V. CONCLUSION AND FUTURE WORK

To the best of our knowledge, this was the first attempt at training a model for the recognition and linking of named entities to a Wikidata knowledge base for the Serbian language, employing spaCy and a CNN network. Although the results on the test set are satisfactory, future research will require expanding the training dataset and applying transformers, which have proven to be more successful than CNN networks in named entity recognition.

One issue observed was the model's difficulty in properly classifying non-linked entities. Increasing the number of linked entities in the knowledge base and expanding the number of properties for extracting QIDs from Wikidata to have more data in the training set would likely improve performance in this area.

Considering that Leximirka has an exhaustive list of locations, a comparison will be made to verify which countries and other categories from Leximirka do not have a corresponding QID. Conversely, an analysis will be conducted to identify which countries, seas, mountains, or organizations exist in the dataset but are missing from Leximirka.

Synthetic sentences have proven to be a valuable source for capturing inflected forms crucial for the morphologically rich Serbian language. Expanding the dataset to include entities with a low frequency of appearance could further enhance the model's performance.

Given the rapid advancements in NLP, CNNs have become somewhat obsolete. Replacing the CNN base layer with BERT or other transformer models is a promising direction for developing a more accurate and robust model.

ACKNOWLEDGMENT

This research was supported by the Science Fund of the Republic of Serbia, #7276, Text Embeddings - Serbian Language Applications - TESLA.

TABLE IV
EVALUATION RESULTS ON AN INDEPENDENT DATASET.

ID	TP	FP	FN	P	R	F ₁
SRPCNNEL						
novel	29	2	10	0.94	0.74	0.83
newspaper	136	0	84	1.00	0.62	0.76
BASELINE SPACY ENTITY LINKER						
novel	2	0	39	1.00	0.05	0.10
newspaper	10	8	202	0.67	0.05	0.09

REFERENCES

- [1] K. Balog, *Entity-oriented search*. Springer Nature, 2018. <https://doi.org/10.1007/978-3-319-93935-3>.
- [2] W. Shen, Y. Li, Y. Liu, J. Han, J. Wang, and X. Yuan, "Entity Linking Meets Deep Learning: Techniques and Solutions," 2021. <https://doi.org/10.1109/TKDE.2021.3090865>.
- [3] R. Hanslo, "Evaluation of Neural Network Transformer Models for Named-Entity Recognition on Low-Resourced Languages," in *2021 16th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pp. 115–119, 2021. <http://dx.doi.org/10.15439/2021F7>.
- [4] W. Shen, J. Wang, and J. Han, "Entity linking with a knowledge base: Issues, techniques, and solutions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 2, pp. 443–460, 2014. <https://dx.doi.org/10.1109/TKDE.2014.2327028>.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 4171–4186, Association for Computational Linguistics, 2019. <https://doi.org/10.18653/v1/N19-1423>.
- [6] W. Yin, M. Yu, B. Xiang, B. Zhou, and H. Schütze, "Simple Question Answering by Attentive Convolutional Neural Network," in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (Y. Matsumoto and R. Prasad, eds.), (Osaka, Japan), pp. 1746–1756, The COLING 2016 Organizing Committee, 2016. <https://doi.org/10.48550/arXiv.1606.03391>.
- [7] T. Lin, Mausam, and O. Etzioni, "Entity linking at web scale," in *Proceedings of the joint workshop on automatic knowledge base construction and web-scale knowledge extraction (AKBC-WEKEX)*, pp. 84–88, Association for Computational Linguistics, 2012. <https://aclanthology.org/W12-3016>.
- [8] K. Labusch and C. Neudecker, "Named Entity Disambiguation and Linking Historic Newspaper OCR with BERT," in *CLEF (Working Notes)*, p. 33, CEUR-WS, 2020. http://ceur-ws.org/Vol-2696/paper_163.pdf.
- [9] Z. Liu, Y. Leng, M. Wang, and C. Lin, "Named Entity Recognition and Named Entity on Esports Contents," in *2020 15th Conference on Computer Science and Information Systems (FedCSIS)*, pp. 189–192, 2020. <https://doi.org/10.15439/2020F24>.
- [10] X. Liu, Y. Li, H. Wu, M. Zhou, F. Wei, and Y. Lu, "Entity linking for tweets," in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1304–1311, Association for Computational Linguistics, 2013. https://doi.org/10.1142/9789813227927_0019.
- [11] E. French and B. T. McInnes, "An overview of biomedical entity linking throughout the years," *Journal of Biomedical Informatics*, vol. 137, p. 104252, 2023. <https://doi.org/10.1016/j.jbi.2022.104252>.
- [12] R. Sharma, D. Chauhan, and R. Sharma, "Named Entity Recognition System for the Biomedical Domain," in *2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pp. 837–840, 2022. <http://dx.doi.org/10.15439/2022F63>.
- [13] I. Guellil, A. Garcia-Dominguez, P. R. Lewis, S. Hussain, and G. Smith, "Entity linking for English and other languages: a survey," *Knowledge and Information Systems*, pp. 1–52, 2024. <https://doi.org/10.1007/s10115-023-02059-2>.
- [14] M. Neumann, D. King, I. Beltagy, and W. Ammar, "SciSpaCy: Fast and Robust Models for Biomedical Natural Language Processing," in *Proceedings of the 18th BioNLP Workshop and Shared Task*, Association for Computational Linguistics, 2019. <https://doi.org/10.18653/v1/w19-5034>.
- [15] O. Bodenreider, "The unified medical language system (UMLS): integrating biomedical terminology," *Nucleic acids research*, vol. 32, no. suppl_1, pp. D267–D270, 2004. <https://doi.org/10.1093/nar/gkh061>.
- [16] G. O. Consortium, "The Gene Ontology (GO) database and informatics resource," *Nucleic acids research*, vol. 32, no. suppl_1, pp. D258–D261, 2004. <https://doi.org/10.1093/nar/gkh036>.
- [17] J. M. Van Hulst, F. Hasibi, K. Dercksen, K. Balog, and A. P. de Vries, "Rel: An entity linker standing on the shoulders of giants," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2197–2200, 2020. <https://doi.org/10.1145/3397271.3401416>.
- [18] N. De Cao, L. Wu, K. Papat, M. Artetxe, N. Goyal, M. Plekhanov, L. Zettlemoyer, N. Cancedda, S. Riedel, and F. Petroni, "Multilingual Autoregressive Entity Linking," *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 274–290, 2022. https://doi.org/10.1162/tacl_a_00460.
- [19] E. Boros, E. L. Pontes, L. A. Cabrera-Diego, A. Hamdi, J. G. Moreno, N. Sidère, and A. Doucet, "Robust named entity recognition and linking on historical multilingual documents," in *Conference and Labs of the Evaluation Forum (CLEF 2020)*, vol. 2696, pp. 1–17, CEUR-WS Working Notes, 2020. <https://doi.org/10.5281/zenodo.4068075>.
- [20] K. Papanitiou, V. Efthymiou, and D. Plexousakis, "Automating Benchmark Generation for Named Entity Recognition and Entity Linking," in *European Semantic Web Conference*, pp. 143–148, Springer, 2023. https://doi.org/10.1007/978-3-031-43458-7_27.
- [21] M. Plekhanov, N. Kassner, K. Papat, L. Martin, S. Merello, B. Kozlovskii, F. A. Dreyer, and N. Cancedda, "Multilingual End to End Entity Linking," *arXiv*, 2023. <https://doi.org/10.48550/arXiv.2306.08896>.
- [22] J. Raiman and O. Raiman, "DeepType: Multilingual Entity Linking by Neural Type System Evolution," 2018. <https://doi.org/10.48550/arXiv.1802.01021>.
- [23] P. Nugues, "Linking Named Entities in Diderot's Encyclopédie to Wikidata," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 10610–10615, 2024. <https://doi.org/10.48550/arXiv.2406.03221>.
- [24] N. Loukachevitch, E. Artemova, T. Batura, P. Braslavski, V. Ivanov, S. Manandhar, A. Pugachev, I. Rozhkov, A. Shelmanov, E. Tutubalina, et al., "NEREL: a Russian information extraction dataset with rich annotation for nested entities, relations, and wikidata entity links," *Language Resources and Evaluation*, pp. 1–37, 2023. <https://doi.org/10.1007/s10579-023-09674-z>.
- [25] A. Delpuch, "Opentapioca: Lightweight entity linking for wikidata," *arXiv preprint arXiv:1904.09131*, 2019. <https://doi.org/10.48550/arXiv.1904.09131>.
- [26] O. Perisic, S. Ranka, I. N. Milica, Š. Mihailo, et al., "It-Sr-NER: CLARIN Compatible NER and Geoparsing Web Services for Italian and Serbian Parallel Text," in *Selected Papers from the CLARIN Annual Conference 2022, Czechia, 2022*, pp. 99–110, Linköping University Electronic Press, 2023. <https://doi.org/10.3384/ecp198010>.
- [27] O. Perišić, S. Ranka, I. N. Milica, and Š. Mihailo, "It-Sr-NER: Web Services for Recognizing and Linking Named Entities in Text and Displaying Them on a Web Map," *Infotheca - Journal for Digital Humanities*, vol. 23, no. 1, pp. 61–77, 2023. <https://doi.org/10.18485/infoteka.2023.23.1.3>.
- [28] Y. Cao, L. Huang, H. Ji, X. Chen, and J. Li, "Bridge Text and Knowledge by Learning Multi-Prototype Entity Mention Embedding," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1623–1633, Association for Computational Linguistics, 2017. <https://doi.org/10.18653/v1/P17-1149>.
- [29] M. Francis-Landau, G. Durrett, and D. Klein, "Capturing Semantic Similarity for Entity Linking with Convolutional Neural Networks," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (K. Knight, A. Nenkova, and O. Rambow, eds.), (San Diego, California), pp. 1256–1261, Association for Computational Linguistics, 2016. <https://doi.org/10.18653/v1/N16-1150>.
- [30] Y. Shi, R. Yang, C. Yin, Y. Lu, Y. Yang, and Y. Tao, "Entity Linking Method for Chinese Short Texts with Multiple Embedded Representations," *Electronics*, vol. 12, no. 12, 2023. <https://doi.org/10.3390/electronics12122692>.
- [31] R. Pozzi, R. Rubini, C. Bernasconi, and M. Palmonari, "Named Entity Recognition and Linking for Entity Extraction from Italian Civil Judgements," in *International Conference of the Italian Association for Artificial Intelligence*, pp. 187–201, Springer, 2023. https://doi.org/10.1007/978-3-031-47546-7_13.
- [32] S. MORAKIS, F. HASIBI, and M. LARSON, "Entity Linking for Greek," 2021.
- [33] R. Stanković, C. Krstev, B. Š. Todorović, and M. Škorić, "Annotation of the Serbian ELTeC Collection," *Infotheca—Journal for Digital Humanities*, vol. 21, no. 2, pp. 43–59, 2021. <https://doi.org/10.18485/infoteka.2021.21.2.3>.
- [34] D. Vrandečić and M. Krötzsch, "Wikidata: a free collaborative knowledgebase," *Communications of the ACM*, vol. 57, no. 10, pp. 78–85, 2014. <https://doi.org/10.1145/2629489>.
- [35] B. Lazić and M. Škorić, "From DELA based dictionary to Leximirka lexical database," *Infotheca—Journal for Digital Humanities*, vol. 19, no. 2, pp. 00–00, 2019. <https://doi.org/10.18485/infoteka.2019.19.2.4>.

- [36] D. Hernández, A. Hogan, C. Riveros, C. Rojas, and E. Zerega, "Querying Wikidata: Comparing SPARQL, Relational and Graph Databases," in *The Semantic Web—ISWC 2016: 15th International Semantic Web Conference, Kobe, Japan, October 17–21, 2016, Proceedings, Part II 15*, pp. 88–103, Springer, 2016. https://doi.org/10.1007/978-3-319-46547-0_10.
- [37] D. Vitas, S. Koeva, C. Krstev, and I. Obradović, "Tour du monde through the dictionaries," in *Actes du 27eme Colloque International sur le Lexique et la Grammaire*, pp. 249–256, 2008.
- [38] C. Krstev, D. Vitas, and A. Trtovac, "Orwells 1984—the Case of Serbian Revisited," in *Proc. of 5th Language & Technology Conference*, pp. 25–27, 2011.
- [39] R. Stanković, C. Krstev, D. Vitas, N. Vulović, and O. Kitanović, "Keyword-based search on bilingual digital libraries," in *Semantic Keyword-Based Search on Structured Data Sources: COST Action IC1302 Second International KEYSTONE Conference, IKC 2016, Cluj-Napoca, Romania, September 8–9, 2016, Revised Selected Papers*, pp. 112–123, Springer, 2017. https://doi.org/10.1007/978-3-319-53640-8_10.
- [40] M. Ikonić Nešić, S. Petalinkar, S. Ranka, and Š. Mihailo, "BERT downstream task analysis: Named Entity Recognition in Serbian," in *14th International Conference on Information Society and Technology – ICIST 2024*, unpublished, 2024.
- [41] M. Škorić, "Novi jezički modeli za srpski jezik," *Infotheca - Journal for Digital Humanities*, 2024. <https://doi.org/10.48550/arXiv.2402.14379>.
- [42] Y. Liu, M. Ott, N. Goyal, *et al.*, "Roberta: A robustly optimized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019. <https://doi.org/10.48550/arXiv.1907.11692>.
- [43] J.-C. Klie, M. Bugert, B. Boullosa, *et al.*, "The INCEPTION Platform: Machine-Assisted and Knowledge-Oriented Interactive Annotation," in *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pp. 5–9, 2018.
- [44] B. Šandrih Todorović, C. Krstev, R. Stanković, and M. Ikonić Nešić, "Serbian NER& Beyond: The Archaic and the Modern Intertwined," in *Deep Learning Natural Language Processing Methods and Applications – Proceedings of the International Conference Recent Advances in Natural Language Processing (RANLP 2021)* (G. Angelova, M. Kunilovskaya, R. Mitkov, and I. Nikolova-Koleva, eds.), pp. 1252–1260, INCOMA Ltd., September 2021. https://doi.org/10.26615/978-954-452-072-4_141.
- [45] R. Stanković, C. Krstev, B. Šandrih Todorović, and M. Škorić, "Annotation of the Serbian ELTeC Collection," *Infotheca - Journal for Digital Humanities*, vol. 21, no. 2, pp. 43–59, 2021. <https://doi.org/10.18485/infotheca.2021.21.2.3>.
- [46] B. Lazić and M. Škorić, "From DELA based dictionary to Leximirka lexical database," *Infotheca - Journal for Digital Humanities*, vol. 19, no. 2, pp. 81–98, 2020. <https://10.18485/infotheca.2019.19.2.4>.

Towards crop traits estimation from hyperspectral data: evaluation of neural network models trained with real multi-site data or synthetic RTM simulations

Lorenzo Parigi
0000-0003-4641-7672
Institute for Electromagnetic Sensing of the Environment,
National Research Council, 20133
Milan, Italy.
Department of Civil,
Constructional and Environmental
Engineering, Sapienza University
of Rome, Rome, Italy
Email: parigi.l@irea.cnr.it

Gabriele Candiani
0000-0003-0575-068X
Institute for Electromagnetic Sensing of the Environment,
National Research Council, 20133
Milan, Italy
Email: candiani.g@irea.cnr.it

Ignazio Gallo
0000-0002-7076-8328
Department of Theoretical and Applied Science, University of Insubria, 21100 Varese, Italy
ignazio.gallo@uninsubria.it

Piero Toscano
0000-0001-9184-0707
Institute of BioEconomy, National Research Council,
50145 Florence, Italy
piero.toscano@ibe.cnr.it

Mirco Boschetti
0000-0003-2156-4166
Institute for Electromagnetic Sensing of the Environment, National Research Council, 20133
Milan, Italy
boschetti.m@irea.cnr.it

Abstract—Hyperspectral images from newly launched (ASI-PRISMA and DLR-EnMAP) and future satellite (ESA-CHIME) are an opportunity, thanks to the high spectral resolution and full range continuity, to improve the retrieval of information about the crop parameters and status. The high dimensionality of hyperspectral data and the non-linear relationship between the crop biophysical parameters and their spectral signature make quantitative estimation of crop characteristics challenging, to address these problems we tested different configurations of neural networks (fully connected and convolutional). We tested the different architectures on two training dataset, one consists in ground data collected in three experiments, in different locations and seasons, the second one (hybrid) is composed by synthetic data generated using a radiative transfer model (PROSAIL-PRO). Preliminary results for LAI, CCC and CNC retrieval are encouraging in particular when ground data are exploited demonstrating of the potentiality of NN to fully exploit the information density of the hyperspectral data.

Index Terms—Hyperspectral, Neural Network, RTM, PROSAIL, synthetic data, Wheat

I. INTRODUCTION

IN THE pre-industrial period, worldwide production and consumption of food happened parallel to each other. Nowadays, the global megatrends (climate change, popula-

tion growth, technological change) gradually caused the supply-demand balance to shift towards a not sufficient and unsustainable food production, with a potentially dramatic consequence for environmental and humanitarian aspects [1]. Considering this scenario, food-production system are forced to increase yields while protecting their most important production factors, soil from degradation, water and air from pollution and atmosphere from emissions of greenhouse gasses [2] as well as contributing to the mitigation of climate changes by increasing the carbon stock capability of agroecosystems.

To achieve such a goal, agriculture in the recent decades has increased the interest in collecting complex information about field status to better manage nutrients, water, chemicals [3] and tillage operations. For this scope geo-information products can be used to provide farmers with decision-supporting spatial information (crop traits maps), able to highlight within-field crop variability, as a fundamental tool to support site-specific management (i.e. precision farming) [1], [4]. Among the agro-practices supported by precision farming, nitrogen (N) fertilizations are fundamental, being nitrogen the most important limiting factor to crop growth together with water deficit. Overall, the nitrogen use efficiency in agriculture is estimated at 60%, with a negative effect on the sustainability of crop production from an economic and ecological point of view [5]. Excess nitrogen can be transformed into N_2O , which is a major contributor to climate change; in addition, it can be leached and reach water

The study was performed under the funding of the PNRR MUR – M4C2 (Mission 4 Component 2) Investment 1.4 “National Research Centre for Agricultural Technologies” Agritech, with project code CUP HUB – B63D21015240004.

masses, thus promoting eutrophication [6]. In order to rationalize the use of fertilizers in a smart agriculture paradigm, a deeper knowledge of the spatio-temporal dynamics of crops is required. Monitoring of this variability can be achieved through the use of data acquired from satellite earth observation systems able to produce maps of crop status at regular time.

In this framework, a great contribution is expected by hyperspectral remote sensing data able to provide continuous information to identify specific spectral features diagnostic of soil-plant compounds. New spaceborne image spectroscopy missions have been recently launched e.g. the PRISMA satellite, of the Italian Space Agency (ASI), and EnMAP of the German Space Agency (DLR) are already in orbit from 2019 and 2022 respectively. These systems are precursors able to provide contiguous, spectral sampling, from visible to shortwave infrared regions that covers ranges unobserved by multispectral data (400-2500 nm) and they are used by the scientific community to develop algorithm solutions for the future operational missions such as the CHIME (Copernicus Hyperspectral Imaging Mission for the Environment) of the European Space Agency's (ESA) [7].

A. From spectra to crop status information a complex task

Methods for quantitatively estimating the biophysical parameters (biopars) of plants from hyperspectral data are therefore required in order to generate space-time variability information to support agro-monitoring and management. Different biopars influence specific regions of the electromagnetic spectrum (e.g. Chlorophyll in the visible red-edge region 600-750 nm), but unfortunately several biopars can have a common influence on the same portion of the spectrum. For example, water, nitrogen, and carbon base constituent of leaves strongly influence the short wave infrared region around 1700 nm. Because of this, biopars retrieval from spectral data is not an easy task and classical statistical approaches can fail to provide reliable quantitative estimation, because the relationships are complex and very often non-linear [8], [9].

To overcome the problem of non-linearity and the interaction of the effect of different biopars on the spectrum, various machine learning algorithms have been proposed in the literature. In the field of machine learning, neural networks (NN) have achieved great popularity in recent years, they can be very effective in solving complex problems in both computer vision in agriculture [10] and quantitative estimation of environmental variable because they are able to achieve very good results even in the presence of non-linear relationships [11], [12]. Specifically, we focus on two NN architectures: fully connected networks (FC-NNs) and convolutional neural networks (CNNs). FC-NNs are straightforward to implement and suitable for tabular data. At the same time, CNNs, commonly used in image processing, can effectively capture local spectral features by applying convolutional filters across the spectral bands.

Another issue for the development of retrieval models is in general the small amount of actual ground vs spectral data on which the models should be trained. This scarcity is mainly due to the fact that it is expensive in terms of money and time to perform coupled sampling of biopars and spectroradiometric measurements from ground or contemporary to satellite overpass. Furthermore, the sampled data are representative of only the environment and conditions in the area, thus making it more difficult to export the results obtained with them to other contexts. To solve this problem, the use of methods that exploit physically based models is emerging in the literature; these radiative transfer models (RTMs) allow to simulate spectra from the leaf components and their arrangement in space. The advantage of RTMs is that they can simulate canopy spectral response even for diverse conditions that would be difficult to sample in the field. Once generated such a database of vegetation parameters (model input) and synthetic spectra (model output), machine learning algorithms can be trained to solve the problem. This method in literature is called hybrid and has been recently proposed in scientific literature as the state-of-the-art approach, as it combines the generic properties of RTMs with the flexibility and computational efficiency of machine learning regression algorithms, representing an innovative solution to the so-called inversion problem [11], [13], [14], [15], [16].

In this framework, the aim of this work is to test a machine learning solution to generate a retrieval model able to estimate crop biopar from hyperspectral data. More specifically the study wanted to generate a multi year/site database exploiting ground hyperspectral measurements acquired in different experiments together with corresponding bipolar measurements. This dataset was analyzed with different NN solutions in a ground data driven (using only experimental data) and hybrid approach (using RTM simulation). Different algorithm configurations have been tested, fully-connected (FC) and convolutional neural networks (CNN) to determine which architecture is the most effective. The long term perspective is to develop solutions for automatic estimation of quantitative crop information from new generation satellite data such as PRISMA and EnMAP.

II. MATERIALS AND METHODS

A. Ground data

The experimental data used to test NN solutions have been acquired in the field in 2022 and 2023 from three different projects in 8 farms in Italy (Fig. 1).

We exploited data acquired with different purposes. A Field Phenotyping experiment (FP) was conducted in Arborea (Sardinia) in 2022. 4 different durum wheat varieties were cultivated in experimental plots of 6 m × 3 m with two types of soil preparation and 4 fertilization levels with three replicate for a total of 96 samples. Data were acquired three times during the season. Details of experiment and data acquisition can be found in [17].



Fig. 1: Position of the experiments

Field level Experiments (FE), devoted to set up a decision support system to support sowing density and fertilization management, were performed on 5 farms cultivated with soft wheat in the center-north of Italy in 2023. In each farm, one field was divided in 4 strips: the first representing the control (standard management following farm prescription), the second involves a reduction of nitrogen fertilization (-20%), the third reduction of seed density (-20%) and the last the combination of nitrogen and seeds reduction (-20% of both). Fields were monitored two times during the season at stem elongation and heading.

The Farm Monitoring experiment (FM) was conducted in Jolanda di Savoia (Ferrara) in Bonifiche Ferraresi estate in 2023. Wheat fields conditions were monitored acquiring data on four (10 by 10 m) plots in 10 elementary sampling units (ESU) for a total of 40 sampling sites. ESU were selected in different crop conditions defined according to analysis of previous year satellite data (within field anomaly from Sentinel 2 data) and on the base of the soil maps provided by the farm. 5 ESU were positioned in two fields cultivated with Durum and Soft Wheat respectively.

In all the farms, Leaf Area Index (LAI), Canopy Chlorophyll Content (CCC), and Canopy Nitrogen Content (CNC) were collected following [16] approach. LAI, which measures the one-sided green leaf area per unit of ground surface, is a crucial canopy trait as it describes vegetation density and regulates carbon, energy, and water fluxes in terrestrial ecosystems. Quantifying CCC is essential for monitoring photosynthetic efficiency and the early detection of crop

stress, such as chlorosis. CNC is a direct estimation of plant nitrogen uptake fundamental to support smart fertilization strategy.

Proximal spectral measurements were acquired using a handheld spectrometer (Spectral Evolution) with full spectral range (350-2500 nm) capacity and 3 nm spectral sampling interval (SSI). Spectral measurements were conducted per each plot approximately 1 m above the canopy with a nadir viewing angle and under clear sky conditions with multiple replicates. Before the target measurements, the radiance of a reference panel was collected to derive the top of canopy reflectance for each replicate, then the plot average value was calculated. The ground spectra were resampled to match the PRISMA configuration, encompassing a spectral range from 400 to 2500 nm with an average SSI of less than 10 nm. This resampling process resulted in a total of 230 spectral bands, utilizing Gaussian spectral response functions (SRF) generated based on the center wavelengths and full-width-half-maximum values of PRISMA bands. Spectral regions between 1328-1491 nm and 1794-1993 nm and the last portion of the SWIR between 2378-2500 nm were excluded resulting in a total of 170 so-called PRISMA-like spectral bands. The final goal is to transfer the retrieval models to actual PRISMA images to produce biopars' maps to investigate crop status conditions in space and time.

The ground data from the three experiments is combined in a single database composed by the biopars' values and the corresponding measured reflectance for the 170 PRISMA-like spectral bands. We split all the data in train, validation and test set, our objective was to generate the most representative validation and test sets by assuring, if possible, that all the combinations produced in the different experiments were present in both. The number of biopar-spectra couples is shown in I.

B. Simulated data from RTM

The initial step of a hybrid approach, which involves training a machine learning algorithm with synthetic data, is therefore the generation of the Look Up Table (LUT) of synthetic spectra and corresponding input parameters. The PROSAIL model has been widely used to obtain plant biochemical and structural variables in the agricultural context [18]. The hybrid approach combines the PROSPECT leaf model [19] and the 4SAIL canopy model [20]. The PROSPECT-PRO is the latest version of PROSPECT which was introduced to differentiate specific absorption coefficients of carbon-based (CBC) and protein (CP) leaf constituents. The 4SAIL model requires as input the leaf reflectance and transmittance generated from PROSPECT as well as canopy density information (i.e., LAI), leaf orientation (average leaf angle, ALA), background spectral properties (i.e. soil reflectance) and the illumination and viewing angles (i.e. sun-sensor-target geometry). The resulting combined model (PROSAIL-PRO [21]) is able to simulate reflectance spectra at canopy level as it would be recorded by a remote sensor (i.e. satellite system).




Fig. 2: Distribution of the parameters LAI [$\text{m}^2 \text{m}^{-2}$], CCC [g m^{-2}], and CNC [g m^{-2}] (from left to right) and their split in training (red), validation (green) and test (blue) sets.

The simulation was performed using a MATLAB script of [15] which i) couples the two models and ii) generates distributions of input variables. The generation of the LUT represents a critical step, as it should be representative of vegetation reflectance spectra, including a-priori information on the ranges and distribution of the input variables [22]. To prevent unrealistic combinations of these variables, the method exploits covariances between certain vegetation traits acquired during several field campaigns. The probability density functions (PDFs) and the related ranges used in the simulation have been selected according to actual measured values. The remaining values were selected according to literature or authors' experience as reported in [15]. The script was used to simulate 50000 reflectance spectra from plausible combinations of crop traits. As shown in Fig. 2, the distribution of the parameters used to generate the reflectance spectra is larger than the sampled records, allowing us to train the model on a more extensive and complete dataset with combination that could not be sampled in our working zones. A subsamples based on different LAI levels is shown in Fig. 3.

C. Neural networks

In this study, we employed two neural network architectures: fully connected neural networks (FC-NNs) and convolutional neural networks (CNNs). The choice of these architectures was driven by their distinct capabilities in handling the characteristics of hyperspectral data:

- FC-NNs are a fundamental type of neural network where each neuron is connected to every other neuron in the subsequent layer. This architecture is straightforward and efficient for handling structured data, such as the resampled spectral measurements we obtained. FC-NNs can capture complex, global relationships between input features (spectral bands) and output variables (crop parameters). In our implementation, this architecture is composed by a sequence of FC blocks (FCb), where each FCb is composed by a linear (or fully connected) layer followed by a ReLU activation function and a dropout layer.
- CNNs are particularly effective in processing data with spatial hierarchies, making them well-suited for image and spectral data. By applying convolutional filters, CNNs can capture local patterns and

Table I: Number of biopar-spectra couple. *FF = Field plot level Phenotyping experiment, FE = Field level Experiment, FC =Farm level Monitoring experiment. ** DW = Durum Wheat, SW = Soft Wheat

Study*	Crop**	Training			Validation			Test		
		LAI	CCC	CNC	LAI	CCC	CNC	LAI	CCC	CNC
FP	DW	141	36	27	69	18	12	66	18	12
FM	SW	59	59	59	24	24	24	20	20	20
FE	DW and SW	36	36	35	14	14	14	14	14	11
Total		236	131	121	107	56	50	100	52	43

Table II: Ranges or steps of the hyperparameters

Hyperparameter	Range
Batch size	16, 32, 64, 128
Learning rate	0.1-0.0001
N. of FCb	1-5
N. of CNNb	0-4
Dropout in the FCb	0-0.5 (by 0.05)
Dropout in the CNNb	0-0.5 (by 0.05)
N. of output neurons of the FCb	16, 32, 64, 128, 256, 512
N. of output channels of the CNNb	4, 8, 16, 32
Kernel size of the CNNb	3, 5, 7

spectral features across bands, which is crucial for identifying subtle variations in hyperspectral data that correspond to different crop traits. The hierarchical feature extraction in CNNs enables a more nuanced understanding of the spectral signatures, potentially leading to better performance in crop parameter estimation. In our implementation, the CNN is a sequence of Convolutional blocks (CNNb) and FCb, where each CNNb is composed by two 1D convolutional layer both followed by a ReLU function and stacked together, followed by a dropout layer and a Max Pooling layer.

We tested 100 different models with different hyperparameterization. The hyperparameters that change between each model are: batch size, learning rate, number of FCb, number of CNNb, dropout in the FCb (d) and in the CNNb




Fig. 3: Examples of synthetic hyperspectral signature. The line represent the reflectance mean and the semi-transparent zone is the standard deviation

(d), number of neurons in the FCb (n_{out}), number of output channel in the convolutional layers (c_{out}), and the kernel size of the convolutional layers (k). The ranges of those hyperparameters are described in the II and in Fig. 4 is showed the blocks' architecture. The shape of input data is $[B, 170]$ for the FC architecture and $[B, 1, 170]$ for the CNN architecture, where B is the batch size. The stride (s) of the CNN layers is always 1, the padding value (p) depends on the k value and it is obtained by the integer division of k by 2. The k and s values of the max pool layer is 2 for both.

Given this hyperparameterization strategy, 19 out of 100 models present only FC layers, the remaining present at least one CNNb.

The number of epochs in the training phase was set to 3000 for data driven FC NN, 1000 for the data driven CNN, for the hybrid models 300 and 100 for FC NN and CNN, respectively. The models with the best Mean Squared Error (MSE) on the validation set were selected during the training without reaching the maximum epoch number, then they were used to predict on the test data.

Following the training phase, the top-performing models were assessed on the test dataset. The correlation between observed and estimated data was evaluated using the coefficient of determination (R^2) of linear regression, while the retrieval error was quantified by calculating the Root Mean Squared Error (RMSE) and the relative RMSE (rRMSE) as follows:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - y_i)^2} \quad (1)$$

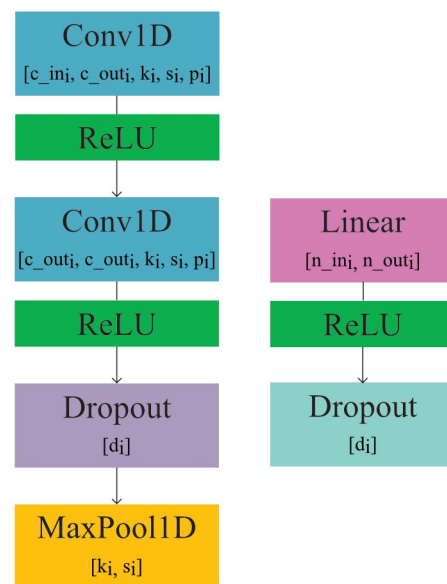


Fig. 4: Description of the convolutional block (CNNb) on the left and the fully connected block (FCb) on the right. Where c_{in} and c_{out} are the input and output size of the CNNs, k , s and p are the kernel, stride and padding size, respectively. n_{in} and n_{out} are the FC dimensions. d is the dropout value. The subscript i is the model number.

$$rRMSE = \frac{RMSE}{y_{max} - y_{min}} \quad (2)$$

Where x_i and y_i are respectively the i -th sample of the predicted and actual values, y_{max} and y_{min} are the max and min value of the observed parameters, respectively.

III. RESULTS AND DISCUSSION

Results for the best ten models, for the three considered biopar (LAI, CCC and CNC) are reported in III, IV, and V together with details of adopted hyperparametrization. Fig. 5, Fig. 6, and Fig. 7 shows the scatter plot of actual vs predicted biopar (LAI, CCC and CNC) for the first five best models reporting regression line for the different dataset (FP, FM, FE).

Models for LAI retrieval (Fig. 5 and III) produced very satisfying results with a R^2 , between observed and predicted values, up to 0.73 and a $rRMSE$ below 12%. Regarding the tested retrieval methods, the data-driven approach ($RMSE = 0.962$) outperforms the hybrid approach ($RMSE = 1.209$) on average, demonstrating a 25% reduction in retrieval error.

The primary distinction between the hybrid and data-driven approaches in the LAI estimation lies in the differing numbers of models that incorporate at least one CNNb. Among the top 10 data-driven models, only three contain one or two CNNb. Furthermore, the first three models are all FC NNs and among them the first two have only one FC layer with 256 and 64 neurons, respectively. Conversely, for the hybrid 7 models with at least one CNNb were selected for the top 10. This aspect may be attributed to the quantity of data employed during the training phase. More complex models, incorporating multiple CNNb and/or deeper architecture, demonstrate superior performance in hybrid approach. This is likely because these model can extract most diagnostic features when trained with 50,000 samples by PROSAIL simulation.

The models for the estimation of CCC (Fig. 6 and IV) produced moderate results in term or correlation between estimates and observation ($R^2 < 0.5$). Once again, the results are superior when the model are trained on field data, with an average $RMSE$ of 0.616 ($rRMSE < 15\%$). However, the scatterplots of Fig. 6 show discrepancies between the experiments, with the FE experiment poorly predicted. On the other hand, in the hybrid approach the FM and FE experiments tend to be more in accordance, instead the FP deviates the most from the 1:1 line. Due to the larger CCC range of FP, the average value of R^2 for hybrid scenario (0.259) is lower than the data driven one (0.471). However, the difference in term of $RMSE$ is only of 13.6% between the best models of the two approaches. Despite the dispersion in the hybrid models, the values tend to cluster by experiment, resulting in relatively tight groupings around their respective trend lines highlighting a possible bias in the ground mea-

surements. In the CCC estimation, both approaches have 5 FC NN models and 5 CNN models in the top 10 best.

The performance of the CNC models (Fig. 7 and V) is similar to the CCC one; however, the values are more scattered, leading to moderate model performance ($R^2 < 0.54$). Additionally, no single experiment consistently outperforms the others. Moreover, the CNC performed slightly worse than CCC in terms of the average $rRMSE$ for both data driven (15.62% vs 17.98%) and hybrid (17.74% vs 18.93%) approaches. The data driven approach produced again better results in term of $RMSE$ (3.270) relatively to the hybrid (3.442), but with a narrower difference of only 5.3%. Despite this negative aspects, it is encouraging that model retrieval are always in the range of the observed values even when hybrid model is used. Moreover, an $rRMSE$ below 20% is an encouraging result, considering the uncertainty of ground measurements in real farming conditions. Similar to the LAI results, 9 of the top 10 hybrid models contain at least one CNNb, compared to only 4 in the data-driven approach. In comparison to the other biopars, LAI exhibits the most significant difference between data driven and hybrid approaches in term of $RMSE$. This discrepancy may be attributed to the larger number of samples on which training have been performed, which enhances the model's capacity to predict field data. Conversely, the estimates of CCC and CNC exhibit a smaller decline in performance when utilizing the hybrid approach. This behavior may be caused by the reduced number of training samples, which may have negatively impacted the performance of the data driven models but candidate hybrid approach as a solution when few data are available. Indeed, the $rRMSE$ of the hybrid models present small variations between LAI, CCC and CNN (in order: 15.65%, 17.74%, 18.93%), underlying a good stability across the biopars estimation that can be used as a starting point when the field data are lacking.

The results are overall encouraging and there are margins of improvements in both data driven and hybrid approaches. The hybrid models tend to perform worse than the data driven ones, but they are promising, especially when the number of samples are relatively small. In this case the generation of spectra through an RTM allows the training of models with good predictive performance on field data that are totally independent hence candidate this model to be more exportable in different situation and data input. Among the biopars analyzed, CNC is the one that has an important agronomic relevance, because the possibility of creating spatio-temporal estimation of plant nitrogen from satellite data is a determining factor in supporting site specific nitrogen fertilization scheduling by producing digital prescription maps. Nitrogen sampling and measurements in field is more labor-intensive than LAI and CCC, consequently the fact that CNC can be estimated using radiometric tools (on ground or from remote) and hybrid models can represent a feasible solution to facilitate smart and more rational crop fertilization dosing.

Table III: Hyperparametrization of the 10 best data driven and hybrid models for the estimation of LAI and their predictive performance. Results are in descending order from best to worst in term of RMSE. CNN model are highlighted in grey

Data driven							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
0	1	[0]	[256]	0.923	0.725	11.94%	
0	1	[0]	[64]	0.933	0.696	12.07%	
0	4	[0]	[512,32,64,16]	0.959	0.721	12.41%	
1	2	[4]	[32,128]	0.962	0.681	12.45%	
0	4	[0]	[64,64,256,32]	0.964	0.703	12.47%	
0	3	[0]	[32,32,256]	0.966	0.678	12.50%	
0	1	[0]	[16]	0.967	0.691	12.51%	
0	4	[0]	[16,128,128,32]	0.98	0.662	12.68%	
1	3	[4]	[16,16,16]	0.981	0.673	12.69%	
2	3	[16,4]	[32,32,256]	0.984	0.671	12.73%	
Mean				0.962	0.690	12.44%	

Hybrid							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
1	2	[4]	[64,512]	1.144	0.519	14.80%	
1	1	[16]	[512]	1.168	0.497	15.11%	
4	2	[4,16,16,8]	[64,256]	1.182	0.509	15.29%	
2	3	[8,8]	[256,16,64]	1.201	0.510	15.54%	
1	4	[4]	[512,256,128,512]	1.216	0.460	15.73%	
0	3	[0]	[32,32,256]	1.225	0.454	15.85%	
1	5	[32]	[64,64,32,256,256]	1.235	0.439	15.98%	
2	2	[32,16]	[128,32]	1.238	0.463	16.02%	
0	5	[0]	[64,512,256,512,512]	1.242	0.441	16.07%	
0	5	[0]	[64,32,256,16,512]	1.243	0.484	16.08%	
Mean				1.209	0.478	15.65%	




Fig. 5: Scatterplot between actual and predicted values of LAI of the 5 best data driven (top) and hybrid (bottom) models. The color represent the experiment FP (Orange), FM (Red), and FE (Blue). The 1:1 line is black, the other are the tendency lines colored by experiment. The FC NN models are represented by triangles and CNN ones by circles

Table IV: Hyperparametrization of the 10 best data driven and hybrid models for the estimation of CCC and their predictive performance. Results are in descending order from best to worst in term of RMSE. CNN model are highlighted in grey

Data driven							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
0	5	[0]	[64,512,256,512,512]	0.616	0.470	14.66%	
0	3	[0]	[256,64,64]	0.623	0.468	14.83%	
0	3	[0]	[32,32,256]	0.638	0.514	15.19%	
1	4	[32]	[256,16,128,64]	0.639	0.488	15.21%	
1	2	[16]	[512,64]	0.648	0.437	15.43%	
0	4	[0]	[16,512,512,512]	0.664	0.405	15.81%	
2	2	[8,32]	[512,32]	0.671	0.489	15.97%	
0	5	[0]	[128,256,512,256,64]	0.684	0.503	16.28%	
1	2	[16]	[512,64]	0.687	0.501	16.35%	
1	2	[16]	[256,256]	0.69	0.440	16.43%	
Mean				0.656	0.471	15.62%	

Hybrid							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
1	2	[16]	[512,64]	0.711	0.298	16.93%	
0	5	[0]	[64,512,256,512,512]	0.725	0.264	17.26%	
0	4	[0]	[16,512,512,512]	0.728	0.287	17.33%	
0	3	[0]	[32,32,256]	0.74	0.261	17.62%	
0	5	[0]	[64,32,256,16,512]	0.747	0.261	17.78%	
2	5	[16,8]	[256,64,512,16,128]	0.749	0.235	17.83%	
2	4	[8,16]	[512,128,128,128]	0.756	0.273	18.00%	
0	5	[0]	[128,32,512,16,128]	0.759	0.218	18.07%	
1	2	[16]	[128,16]	0.767	0.249	18.26%	
1	2	[16]	[512,64]	0.769	0.239	18.31%	
Mean				0.745	0.259	17.74%	




Fig. 6: Scatterplot between actual and predicted values of CCC of the 5 best data driven (top) and hybrid (bottom) models. The color represent the experiment FP (Orange), FM (Red), and FE (Blue). The 1:1 line is black, the other are the tendency lines colored by experiment. The FC NN models are represented by triangles and CNN ones by circles

Table V: Hyperparametrization of the 10 best data driven and hybrid models for the estimation of CNC and their predictive performance. Results are in descending order from best to worst in term of RMSE. CNN model are highlighted in grey

Data driven							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
0	5	[0]	[64,32,256,16,512]	2.837	0.542	15.60%	
3	4	[4,8,16]	[16,32,32,16]	3.232	0.364	17.77%	
0	4	[0]	[16,512,512,512]	3.25	0.388	17.87%	
1	5	[8]	[256,32,128,32,16]	3.281	0.421	18.04%	
0	3	[0]	[32,32,256]	3.288	0.467	18.08%	
1	4	[8]	[64,16,32,256]	3.328	0.452	18.30%	
0	1	[0]	[256]	3.348	0.357	18.41%	
0	4	[0]	[16,128,128,32]	3.364	0.492	18.50%	
1	2	[4]	[128,512]	3.385	0.459	18.61%	
0	1	[0]	[16]	3.388	0.443	18.63%	
Mean				3.270	0.438	17.98%	

Hybrid							
CNNb	FCb	CNN channels	FC neurons	RMSE	R ²	rRMSE	
4	5	[32,4,8,32]	[64,32,16,512,128]	3.114	0.366	17.12%	
1	5	[32]	[64,64,32,256,256]	3.216	0.313	17.68%	
1	1	[16]	[16]	3.359	0.346	18.47%	
1	2	[16]	[128,16]	3.409	0.353	18.74%	
2	4	[8,16]	[512,128,128,128]	3.52	0.296	19.35%	
0	5	[0]	[64,512,256,512,512]	3.538	0.262	19.45%	
1	2	[32]	[32,512]	3.543	0.390	19.48%	
2	4	[8,4]	[16,256,512,512]	3.562	0.291	19.58%	
1	4	[4]	[512,256,128,512]	3.576	0.240	19.66%	
2	2	[8,32]	[512,32]	3.587	0.312	19.72%	
Mean				3.442	0.317	18.93%	




Fig. 7: Scatterplot between actual and predicted values of CNC of the 5 best data driven (top) and hybrid (bottom) models. The color represent the experiment FP (Orange), FM (Red), and FE (Blue). The 1:1 line is black, the other are the tendency lines colored by experiment. The FC NN models are represented by triangles and CNN ones by circles

IV. CONCLUSION AND FUTURE PERSPECTIVE

The dataset involved in this work come from different location and was sampled by different operators. Integrating different datasets is challenging, as cardinality is not the only crucial factor; data quality and consistency in data collection and measurement protocols are also fundamental. In future work, the quality of individual datasets must be evaluated, and outliers and anomalous values must be assessed. Additionally, a deeper analysis of plant samples and measurement protocols is necessary. New data from ongoing experiment conducted in 2024 will be tested to assess the exportability of the obtained models on an independent data set and to train new model on increased datasets.

To refine the hybrid model's performance active learning and transfer learning techniques on field data will be explored. Finally, models trained with ground data will be tested on actual satellite data to validate their applicability to spaceborne sensors. The integration of the ground and satellite information, along with RTM simulations will be evaluated for an advanced retrieval scheme.

ACKNOWLEDGMENT

We thanks Dr. Ramin Heidarian Dehkordi for preparing the FP dataset acquired in 2022 in Sardinia and Dr. Riccardo Dainelli for the acquisition of data in the FE condition in 2023.

REFERENCES

- [1] T. B. Hank *et al.*, "Spaceborne Imaging Spectroscopy for Sustainable Agriculture: Contributions and Challenges," *Surv. Geophys.*, vol. 40, no. 3, pp. 515–551, May 2019, doi: 10.1007/s10712-018-9492-0.
- [2] FAO, Ed., *The state of food and agriculture - Climate change, agriculture and food security*. in The state of food and agriculture, no. 2016. Rome: FAO, 2016.
- [3] M. Wójtowicz, A. Wójtowicz, and J. Piekarczyk, "Application of remote sensing methods in agriculture," 2016.
- [4] P. J. Zarco-Tejada, N. Hubbard, and P. Loudjani, "Precision agriculture: an opportunity for EU farmers: potential support with the CAP 2014-2020," 2014.
- [5] L. Lassaletta *et al.*, "Nitrogen use in the global food system: past trends and future trajectories of agronomic performance, pollution, trade, and dietary demand," *Environ. Res. Lett.*, vol. 11, no. 9, p. 095007, Sep. 2016, doi: 10.1088/1748-9326/11/9/095007.
- [6] S. M. Ogle, K. Butterbach-Bahl, L. Cardenas, U. Skiba, and C. Scheer, "From research to policy: optimizing the design of a national monitoring system to mitigate soil nitrous oxide emissions," *Curr. Opin. Environ. Sustain.*, vol. 47, pp. 28–36, Dec. 2020, doi: 10.1016/j.cosust.2020.06.003.
- [7] K. Berger *et al.*, "Crop nitrogen monitoring: Recent progress and principal developments in the context of imaging spectroscopy missions," *Remote Sens. Environ.*, vol. 242, p. 111758, Jun. 2020, doi: 10.1016/j.rse.2020.111758.
- [8] P. J. Curran, "Remote sensing of foliar chemistry," *Remote Sens. Environ.*, vol. 30, no. 3, pp. 271–278, Dec. 1989, doi: 10.1016/0034-4257(89)90069-2.
- [9] Y. Fu *et al.*, "An overview of crop nitrogen status assessment using hyperspectral remote sensing: Current status and perspectives," *Eur. J. Agron.*, vol. 124, p. 126241, Mar. 2021, doi: 10.1016/j.eja.2021.126241.
- [10] G. Castellano, P. D. Marinis, and G. Vessio, "Applying Knowledge Distillation to Improve Weed Mapping With Drones," presented at the 18th Conference on Computer Science and Intelligence Systems, Sep. 2023, pp. 393–400. doi: 10.15439/2023F960.
- [11] J. Verrelst *et al.*, "Quantifying Vegetation Biophysical Variables from Imaging Spectroscopy Data: A Review on Retrieval Methods," *Surv. Geophys.*, vol. 40, no. 3, pp. 589–629, May 2019, doi: 10.1007/s10712-018-9478-y.
- [12] I. Gallo, M. Boschetti, A. U. Rehman, and G. Candiani, "Self-Supervised Convolutional Neural Network Learning in a Hybrid Approach Framework to Estimate Chlorophyll and Nitrogen Content of Maize from Hyperspectral Images," *Remote Sens.*, vol. 15, no. 19, p. 4765, Sep. 2023, doi: 10.3390/rs15194765.
- [13] M. Weiss, F. Jacob, and G. Duveiller, "Remote sensing for agricultural applications: A meta-review," *Remote Sens. Environ.*, vol. 236, p. 111402, Jan. 2020, doi: 10.1016/j.rse.2019.111402.
- [14] K. Berger *et al.*, "Retrieval of aboveground crop nitrogen content with a hybrid machine learning method," *Int. J. Appl. Earth Obs. Geoinformation*, vol. 92, p. 102174, Oct. 2020, doi: 10.1016/j.jag.2020.102174.
- [15] G. Candiani *et al.*, "Evaluation of Hybrid Models to Estimate Chlorophyll and Nitrogen Content of Maize Crops in the Framework of the Future CHIME Mission," *Remote Sens.*, vol. 14, no. 8, p. 1792, Apr. 2022, doi: 10.3390/rs14081792.
- [16] G. Tagliabue *et al.*, "Hybrid retrieval of crop traits from multi-temporal PRISMA hyperspectral imagery," *ISPRS J. Photogramm. Remote Sens.*, vol. 187, pp. 362–377, May 2022, doi: 10.1016/j.isprs.2022.03.014.
- [17] R. Heidarian Dehkordi *et al.*, "Towards an Improved High-Throughput Phenotyping Approach: Utilizing MLRA and Dimensionality Reduction Techniques for Transferring Hyperspectral Proximal-Based Model to Airborne Images," *Remote Sens.*, vol. 16, no. 3, p. 492, Jan. 2024, doi: 10.3390/rs16030492.
- [18] K. Berger, Z. Wang, M. Danner, M. Wocher, W. Mauser, and T. Hank, "Simulation of Spaceborne Hyperspectral Remote Sensing to Assist Crop Nitrogen Content Monitoring in Agricultural Crops," in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, Valencia: IEEE, Jul. 2018, pp. 3801–3804. doi: 10.1109/IGARSS.2018.8518537.
- [19] J.-B. Féret, "PROSPECT-PRO for estimating content of nitrogen-containing leaf proteins and other carbon-based constituents," *Remote Sens. Environ.*, 2021.
- [20] W. Verhoef, "Light Scattering by Leaf Layers with Application to Canopy Reflectance Modeling: The SAIL Model," 1984.
- [21] J.-B. Féret and F. de Boissieu, *prosail: PROSAIL leaf and canopy radiative transfer model and inversion routines*. 2023. [Online]. Available: <https://gitlab.com/jbferet/prosail>
- [22] M. Weiss, S. Jay, and F. Baret, "S2ToolBox Level 2 products: LAI, FAPAR, FCOVER - Version 1.1.," 2016.

AI-Based Spatiotemporal Crop Monitoring by Cloud Removal in Satellite Images

Jiří Pihrt*, Petr Šimánek*, Alexander Kovalenko*, Jiří Kvapil[†], and Karel Charvát[‡]

*Faculty of Information Technology, Czech Technical University in Prague

Thákurova 9, Praha 6, Czech Republic

Email: petr.simanek@fit.cvut.cz

[†]Lesprojekt s.r.o.

Martinov 197, Záryby, Czech Republic

Email: jiri.kvapil@lesprojekt.cz

[‡]HELP SERVICE – REMOTE SENSING s.r.o.

Husova 2117, Benešov, Czech Republic

Email: charvat@hsrs.cz

Abstract—Efficient crop monitoring and crop dynamics forecasting leveraging diverse satellite and point data are described. UnCRtainTS neural network architecture is utilized for cloud removal in satellite imagery which overcomes an issue in crop monitoring. Combining optical (Sentinel-2) and radar (Sentinel-1) satellite data improves the robustness and accuracy of the model in terms of satellite image reconstruction and vegetation index estimation. However, available soil-type geographical data and land surface analysis products, do not improve prediction accuracy significantly.

I. INTRODUCTION

PRECISION agriculture [16] aims to maximize the output from farming by defining the precise and sufficient amounts of inputs like water, fertilizer, pesticides, etc. at the correct time to the crop for increasing its productivity and maximizing its yields. However, this approach increases the vulnerability of crop health, in the case of input errors [3] as the goal of precision agriculture is to use minimally required amount of inputs. Therefore, efficient and timely crop monitoring is essential to detect harmful patterns that may emerge under certain conditions, potentially compromising the harvest. Factors influencing crop development in the early stages are numerous and complex[9], often requiring expert knowledge to interpret and understand their interactions. These intricacies can lead to important relationships being overlooked, emphasizing the need for advanced monitoring techniques.

In recent years, machine learning has achieved significant success in remote sensing imagery [28],[2], particularly in crop detection using multispectral satellite data [29]. This technology allows for high-accuracy crop identification, providing valuable insights into agricultural practices and crop health. However, despite these advancements, cloud cover [29] remains as a significant obstacle to efficient crop monitoring. This issue is particularly relevant in regions like Central Europe, where average yearly cloud coverage can exceed

50%¹. In such conditions, the efficiency of monitoring crops using satellite images is severely compromised, making it nearly impossible to rely solely on crop detection by remote sensing imagery.

Another critical challenge in agriculture is forecasting harvest dynamics [13]. External factors such as temperature, wind, and precipitation affect crop growth in various ways. These environmental variables can negatively impact crop development, which, if not mitigated timely, can lead to significant yield losses. To address these issues, treatments such as adjusting soil fertilization or irrigation can be applied to mitigate the adverse effects [1]. Understanding how crops grow and respond over time to varying environmental conditions allows for the timely implementation of interventions, thereby reducing the risk of damage and optimizing crop health [8].

Moreover, precise detection of crop dynamics in response to external interventions, such as fertilizing, irrigation, or the application of pesticides and fungicides, can lead to more efficient use of these treatments [26]. This precision agriculture approach can result in significant resource savings, such as conserving water and using fertilizers more reasonably. This helps in limiting the use of pesticides and fungicides to the minimum required for effective pest and disease management, promoting more sustainable agricultural practices.

To overcome the limitations posed by cloud cover and to enhance the forecasting of harvest dynamics, it is essential to integrate machine learning techniques with diverse data sources. This integration can develop robust models capable of handling incomplete or obscured data and improving the resilience and effectiveness of agricultural monitoring systems. This paper aims to investigate these methods, offering solutions to improve crop monitoring and forecasting in the face of frequent cloud cover and other environmental challenges. To

¹<https://www.dwd.de>

predict and forecast crop health we monitored the normalized difference vegetation index (NDVI) which is a standard metric for quantifying the health and density of vegetation using spectral satellite data from red and near-infrared bands.

As a solution to the abovementioned problems, we explored the possibility of forecasting crop dynamics in high cloud coverage conditions. We leveraged a combination of Satellite (Sentinel 2 [15]) and radar data (Sentinel 1 [24]) and point data, such as soil type and land surface analysis (LSA) to develop a model for the next frame prediction to provide cloudless prediction that learns crop development pattern dynamics. By leveraging advanced technologies, more sustainable and efficient agricultural practices can be achieved, to secure a more stable food supply.

II. RELATED RESEARCH

Remote sensing using multispectral satellite data has been widely explored for its applications in agriculture. Multispectral imagery allows for the identification of various crop types and their health status based on their spectral signatures. Studies such as those by Peña-Barragán et al. [14] and Quan et al. [17] have demonstrated the effectiveness of using multispectral data for crop classification and harvest health monitoring. Additionally, dash et al. [4] explore the effectiveness of unmanned aerial vehicles (UAV) and satellite imagery for monitoring forest health, specifically focusing on mature *P. radiata* trees. This research under controlled experimental conditions shows that both UAV and satellite sensors can detect plant stress, as evidenced by deviations in spectral indices and strong correlations with field observations.

However, as it is mentioned above the issue of cloud cover remains a significant barrier. Techniques to mitigate this include the use of cloud masking algorithms and temporal interpolation methods. For example, one of the early works by Zhu et al. [30] presents the Fmask algorithm, which effectively identifies and masks clouds in Landsat imagery, enabling clearer analysis of vegetation.

To overcome the limitations of optical data due to cloud cover, researchers have increasingly turned to radar imagery, which can penetrate clouds and provide consistent data. Sentinel-1 synthetic aperture radar (SAR) data, for instance, has been successfully integrated with optical data to improve crop monitoring. The study by Veloso et al. [25] reports the use of Sentinel-1 and Sentinel-2 data fusion for crop-type mapping, demonstrating improved accuracy and reliability.

Later advances of cloud removing algorithms included using generative adversarial networks for filmy cloud removal on satellite imagery with multispectral conditional [7], using a deep residual neural network and SAR-optical data fusion [12], spatiotemporal generative networks [22], enhanced cloud removal with global-local fusion leveraging data from synthetic aperture radar [27], as well as using uncertainty quantification for cloud removal in optical satellite time series [6], which is the current state-of-the-art method to the best of our knowledge.

Regarding crop classification and crop dynamics monitoring, machine learning techniques, particularly convolutional neural networks (CNNs) and recurrent neural networks (RNNs), have shown great promise in enhancing the analysis of remote sensing data. The work by Russwurm and Körner [19] employed RNNs to analyze time-series data from Sentinel-2. Forecasting crop dynamics involves predicting how crops will develop over time, considering various environmental factors. Models such as those developed by Han et al. [8] leverage weather data and soil conditions to predict crop yields, providing valuable insights for proactive agricultural management.

Moreover, hybrid models combining machine learning with mechanistic crop models offer enhanced predictive capabilities. For example, Shahhosseini et al. [23] integrated a machine learning approach with the Agricultural Production Systems sIMulator (APSIM) crop model, resulting in improved predictions of crop growth and yield under varying environmental conditions. The precise detection and analysis of crop responses to interventions such as fertilization and irrigation can lead to more efficient resource use. Studies such as Lobell et al. [11] have shown that remote sensing can effectively monitor crop responses to nitrogen application, optimizing fertilizer use and enhancing sustainability. Recent advancements in IoT and sensor networks further support this precision agriculture approach, as demonstrated by Liakos et al. (2018) [10], who reviewed the integration of IoT in agriculture for real-time monitoring and decision-making.

III. DATA

Integrating data from multiple sources in machine learning provides numerous advantages, particularly in fields that require comprehensive analysis and forecasting. Combining diverse data sets introduces greater variability and richness, which helps machine learning models capture a wider range of patterns and relationships. This approach leads to more robust and generalized models capable of making accurate predictions in varied scenarios. Additionally, using multiple data sources helps mitigate biases present in individual datasets, resulting in a more balanced representation of the underlying phenomena and reducing both bias and variance in the model's predictions.

In the context of crop dynamics forecasting, the integration of data from satellites, radars, and point sources such as temperature, wind, and precipitation provides significant benefits. Satellite imagery offers valuable insights into crop health and development through multispectral data, but its effectiveness is often limited by cloud cover. Radar data, which can penetrate clouds, ensures continuous monitoring regardless of weather conditions. While radar data provides structural information about crops, it may lack the spectral details available in satellite images. By combining these sources, a more complete and reliable view of crop conditions can be achieved.

It is worth noting that there are publicly available datasets designed for the cloud removal task using machine learning, the most widely utilized currently being SEN12MS-CR-

TS [5]. This dataset contains pairs of spatially aligned Sentinel 1 and 2 images. However, the Sentinel 2 images are of the L1C product. For surface analysis, the atmospherically corrected L2A product is more suitable. Another limitation is the temporal resolution of this dataset, consisting of only 30 samples per year, which is lower than the actual acquisition rate. High intervals between frames in the input sequence may cause missing valuable information in the skipped samples, such as recent cloudless images. Furthermore, the overall dataset contains diverse regions globally. However, this work was specifically aimed at agricultural lands in the Czech Republic. For these reasons, we opted to create our own dataset for this work.

The primary inputs include the most recent Sentinel-2 L2A frames, consisting of all 12 spectral bands. Additionally, we incorporate the most recent Sentinel-1 GRD (IW mode) frames, specifically the VV and VH polarization bands, which are orthorectified and terrain-corrected to ensure spatial alignment with other sources. These datasets are essential for capturing both optical and radar imagery, providing a comprehensive view of the agricultural landscape.

Furthermore, we have experimented with integrating various EUMETSAT LSA products. Among these, we used MDIDSSF, MDMETv3, and MNSLF, which are available at much lower spatial resolution than Sentinel images. For these products, we utilized single values nearest to the region of interest to supplement the primary data. The temporal resolution for these inputs is one day, ensuring that we capture daily variations in the land surface conditions.

Additionally, we considered geographical data from the Czech Republic's VUMOP mapping service (<https://mapy.vumop.cz/>), specifically the soil types layer. This layer, available in the EPSG:5514 coordinate system for the Czech Republic, required reprojection to align with our other datasets. Being categorical data, it was necessary to convert it into a one-hot encoded format to be usable in our models. It is important to note that this data is static, and lacking temporal variation, but it provides valuable contextual information about soil types.

For training purposes, we utilized the Scene Classification Layer (SCL) from all Sentinel-2 frames. This layer is useful for masking out clouds and cloud shadows, ensuring target ground truth data is cloudless. Additionally, we experimented with a binary mask to exclude non-agricultural land from the analysis, derived from a layer available on the VUMOP mapping service. This mask helps in focusing the model's attention solely on agricultural areas, potentially improving the relevance and accuracy of the inferences.

Using the data sources described previously, we create a specialized dataset for this work. Regions of interest are sampled uniformly in the bounding box of the Czech Republic. After removing regions with little to no agricultural land (mapped by VUMOP) or missing data, the number of regions is 212. We use data from the year 2022 for training and 2023 for validation. Only data from April to October are used, as winter months are less useful for crop monitoring.

All spatial data inputs were standardized to a 10-meter resolution to maintain consistency across datasets. For datasets with lower resolutions, we upscaled them to meet this standard. The dataset contains and therefore our models were trained using input frames of 256x256 pixels, allowing for efficient processing while maintaining high spatial detail.

IV. EXPERIMENT DESIGN

For a cloud removal task using machine learning, training samples consist of a cloudless Sentinel-2 target frame, and data that the model uses as input, most importantly a sequence of the most recent Sentinel-2 frames. To generate such samples, we generate sequences of Sentinel-2 frames using a sliding window and pick the last frame as the target. The sample is dropped if the target frame contains clouds or cloud shadows, as classified in SCL. Additional input data, such as a sequence of most recent Sentinel-1 frames, are added depending on the experiment. In our experiments, we standardized the number of recent frames to 5, though this parameter is adjustable depending on the specific requirements of future studies.

The neural network architecture we opted to use is UnCRtainTS [6], as it currently is the state-of-the-art method for the cloud removal task with a publicly available implementation. UnCRtainTS is an attention-based convolutional network, as described in Fig. 1. However, in our experiments, we used the mono-temporal version, which excludes the temporal aggregator. The multi-temporal version requires a significant amount of memory during training. Therefore, due to hardware limitations, we used the mono-temporal version. The advantage of this version is that it allows for more straightforward aggregation of additional data which is either static or not temporally aligned with the Sentinel-2 frames. It is worth noting that Sentinel-1 frames are also not temporally aligned with Sentinel-2.

To confirm the effectiveness of UnCRtainTS on our dataset, we also compare it to a baseline U-Net [18], a widely used image-to-image convolutional network.

For both UnCRtain and U-Net, input sequences are flattened and all inputs are concatenated in the channel dimension. In the case of point data, each value first has to be upscaled into a 256x256 image, so that it can be concatenated with the other images.

Although UnCRtainTS can be used with a negative log-likelihood (NLL) loss function for uncertainty estimation as in the original paper, we used classic regression loss functions (MSE, MAE), to focus on more accurate predictions instead of uncertainty quantification. All models in the experiments are trained using the Adam optimizer with a learning rate of 0.001.

V. RESULTS

The experiments test various input sources and hyperparameters. Unless stated otherwise, by default the model is the following: UnCRtainTS architecture, MAE loss function, batch size = 6, MAE loss function, and inputs are Sentinel-1 and Sentinel-2 sequences.

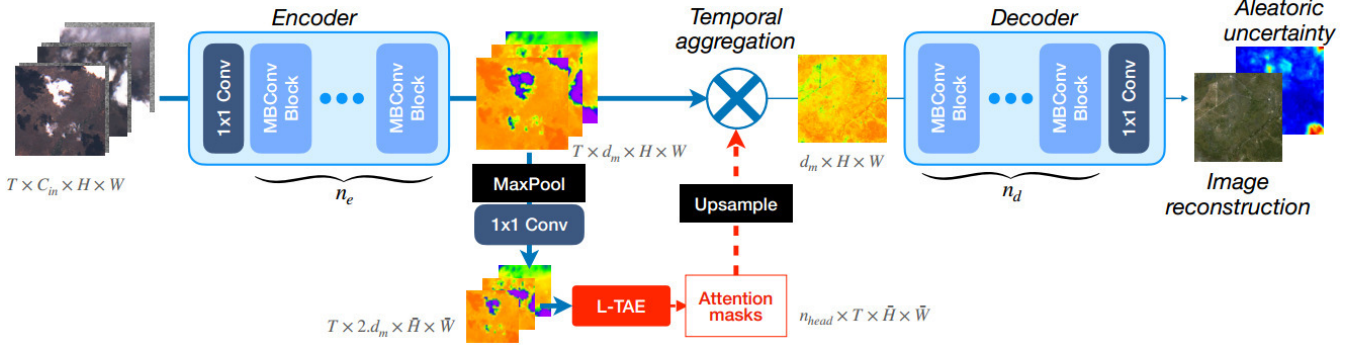


Fig. 1. UnCRtainTS [6] architecture. The network has three main parts; an encoder and decoder consisting of MBConv [21] blocks processing feature maps at full input resolution, and an attention-based temporal aggregator computing an attention mask by applying an L-TAE [20] to downsampled feature maps.

The MAE and MSE metrics, reported in each experiment, are computed on the validation set, exclusively on agricultural land pixels, averaged across all 12 Sentinel-2 L2A bands.

Table I compares UnCRtainTS against several baselines. *Most recent cloudless* simply repeats the most recent fully cloudless Sentinel-2 frame. *Mosaicking* uses the SCL cloud mask to repeat the most recent unobstructed frame for each pixel independently. This uses more recent data but suffers from cloud and cloud shadow artifacts due to imperfect cloud detection, which is mostly not an issue in fully cloudless frames.

The importance of Sentinel-1 and Sentinel-2 input sequences is demonstrated in Table II. It is clear that using both data sources is crucial for good results. Table III experiments with several batch sizes.

Table IV shows that training with MAE loss function outperforms MSE, even on the MSE metric. Additionally, training on all pixels, rather than just those mapped as agricultural land, slightly improves performance. This suggests that incorporating more diverse data makes the model more robust in the target domain (agricultural land) as well. It could also be caused by false positives in the VUMOP mask, and the model trained on all data handling these outliers better.

Experiments including the additional inputs (soil types and LSA point data) are shown in Table V. Interestingly, the inclusion of these additional data sources did not lead to any significant improvement in the model's performance. This suggests that the current model already captures the essential features required for accurate predictions, or that the added data may require further preprocessing or different integration methods to be beneficial.

In the context of crop monitoring, there are metrics useful for quantifying crop health, such as NDVI (normalized difference vegetation index). NDVI is computed from multi-spectral satellite imagery as

$$\text{NDVI} = \frac{\text{NIR} - \text{Red}}{\text{NIR} + \text{Red}}.$$

For Sentinel-2, NIR is band 8 and Red is band 4. Table VI shows that optimizing NDVI directly (optimizing

TABLE I
EVALUATION OF DIFFERENT METHODS

Method	MAE	MSE
most recent cloudless	0.0371	0.00338
mosaicking	0.0386	0.00388
U-Net	0.0278	0.00174
UnCRtainTS	0.0255	0.00153

TABLE II
EVALUATING CONTRIBUTIONS OF SENTINEL INPUT SEQUENCES

Model	Sentinel-1	Sentinel-2	MAE	MSE
UnCRtainTS	✓		0.0375	0.00300
UnCRtainTS		✓	0.0281	0.00191
UnCRtainTS	✓	✓	0.0255	0.00153

$loss(\text{NDVI}(x), \text{NDVI}(y))$ instead of $loss(x, y)$) leads to significantly more accurate NDVI predictions. The disadvantage is that the model no longer predicts the raw Sentinel-2 bands. For applications where that is needed as well, separate models for each task can be trained.

Example predictions using the default model are visualized in Fig. 2. Example NDVI predictions over time using the model trained with NDVI MAE loss function are shown in Fig. 3.

VI. LIMITATIONS

Despite the promising results and potential applications of our crop monitoring system, several limitations must be acknowledged. Cloud coverage and data availability remain significant challenges. Although combining Sentinel-1 and

TABLE III
EVALUATION OF THE MODEL WITH VARYING BATCH SIZE

Model	batch size	MAE	MSE
UnCRtainTS	4	0.0256	0.00159
UnCRtainTS	6	0.0255	0.00153
UnCRtainTS	8	0.0257	0.00156




Fig. 2. Prediction examples. Rows: different samples, first three are from the same region of interest. Columns: visualization of the most recent input Sentinel-1 frame; visible channels of the most recent input Sentinel-2 frame; visible channels of the target and prediction; visualization of NDVI (0-1 maps to yellow-green) of the target and prediction.

TABLE V
EVALUATING VARIOUS LOSS FUNCTIONS (*masked* LOSS IS COMPUTED ONLY ON PIXELS MAPPED AS AGRICULTURAL LAND)

Model	loss function	MAE	MSE
UnCRtainTS	MAE	0.0255	0.00153
UnCRtainTS	MAE masked	0.0258	0.00158
UnCRtainTS	MSE	0.0267	0.00161
UnCRtainTS	MSE masked	0.0291	0.00184

TABLE V
EVALUATING CONTRIBUTIONS OF AUXILIARY INPUTS (VUMOP SOIL TYPES LAYER AND LSA POINT DATA)

Model	soil types	LSA	MAE	MSE
UnCRtainTS			0.0255	0.00153
UnCRtainTS	✓		0.0254	0.00153
UnCRtainTS		✓	0.0258	0.00158
UnCRtainTS	✓	✓	0.0260	0.00160

TABLE VI
EVALUATING NDVI PERFORMANCE WITH VARYING LOSS FUNCTIONS

Model	loss	NDVI MAE	NDVI MSE
UnCRtainTS	NDVI MAE	0.0631	0.0105
UnCRtainTS	NDVI MSE	0.0647	0.0098
UnCRtainTS	MAE	0.0721	0.0122
UnCRtainTS	MAE	0.0801	0.0135

Sentinel-2 data helps mitigate cloud cover issues, there are still instances where data from both sources may be inadequate.

The temporal resolution of available data can also be limited. While we aimed for a daily temporal resolution for some inputs, Sentinel-2 images are often only available every 5 days for specific regions. This gap can lead to temporal inconsistencies and impact the monitoring of rapid changes in crop conditions.

The spatial resolution of some input data required upscaling to match the 10-meter resolution standard. This is particularly




Fig. 3. Examples of NDVI prediction over time. This was acquired by running the model for each sample over the year for a given region and measuring NDVI at three handpicked locations (agricultural land). Target values obscured by clouds are regarded as missing values and not plotted. However, the cloud mask is not perfect and some target values may be affected by clouds.

true for the LSA products that originally had much lower resolutions. Upscaling may introduce artifacts or reduce the precision of the information, which can affect overall model performance.

While we have included optional data inputs such as additional LSA products and specific categorical layers, these have not yet demonstrated significant benefits in our current framework. Future work will continue to explore and validate the potential contributions of these additional datasets to further enhance the accuracy and robustness of the agricultural monitoring models.

The generalizability of our models to other regions, crop types, and environmental conditions is not guaranteed. Differences in agricultural practices, climatic conditions, and soil properties may require additional tuning and validation to ensure the models perform well in diverse settings.

The computational resources required for training and deploying advanced machine learning models, like UnCRtainTS, are substantial. High-performance computing environments or cloud-based solutions are necessary, which may not be accessible to all users, especially in resource-constrained settings. Errors in preprocessing steps, such as cloud masking or data alignment, can propagate through the modeling pipeline, affecting the final predictions' accuracy. Ensuring high-quality preprocessing is critical but can be challenging given the complexity and volume of the data.

Using point data, which includes measurements of environmental factors such as temperature, wind, and precipitation, could be critical for understanding crop dynamics but is not directly observable from remote sensing data. Including these point measurements could provide essential context for interpreting remote sensing data. For example, temperature affects plant metabolism and growth rates, and integrating temperature data helps understand the impact of these factors

on crop development. Precipitation is vital for soil moisture and overall plant health, and combining precipitation data with satellite and radar imagery would help assess drought conditions or waterlogging. Wind influences pollination and the spread of pests and diseases, and wind data can presumably help predict potential pest outbreaks or physical damage to crops. However, these types of data were not available for the Czech Republic.

Additionally, ongoing research and collaboration with agricultural experts can help refine the models, improve data integration techniques, and enhance the system's robustness and applicability across different agricultural contexts.

VII. CONCLUSION

The integration of machine learning techniques with diverse data sources has demonstrated significant potential in enhancing crop monitoring and forecasting, particularly in the context of frequent cloud cover and other environmental challenges. Our study leveraged a combination of the most recent Sentinel-2 L2A frames, Sentinel-1 GRD frames, and various point data sources to develop a robust model for predicting crop dynamics under high cloud coverage conditions. By combining optical and radar imagery, we were able to create a more complete and reliable view of crop conditions, ensuring continuous monitoring and accurate forecasting regardless of weather conditions.

The results of our experiments indicate that the integration of multiple data sources, such as Sentinel-2 L2A frames, and Sentinel-1 GRD frames enhances the accuracy of the models and provides a comprehensive understanding of crop health and development in terms of predicting and forecasting the NDVI index. For instance, the use of radar data from Sentinel-1 complemented the optical data from Sentinel-2 by providing information even under cloud cover.

We also explored the inclusion of various LSA products and categorical data from the VUMOP mapping service, although their contributions to the overall model performance were not as significant. This highlights the importance of selecting relevant and high-quality data sources tailored to specific agricultural monitoring needs.

The next steps for advancing our crop monitoring system involve refining the existing models to improve performance and reduce computational requirements. This includes optimizing model architectures and experimenting with additional data sources, such as temperature, wind, precipitation, and/or soil moisture, as well as sensors and UAV imagery. Extensive field validation will be conducted to assess the accuracy and reliability of the system in real-world conditions, complemented by detailed error analysis to refine the models.

ACKNOWLEDGMENT

This project is funded with state support from the Technology Agency of the Czech Republic and the Ministry of Industry and Trade of the Czech Republic under the TREND Program.

REFERENCES

- [1] Emmanuel Abiodun Abioye et al. “Precision irrigation management using machine learning and digital farming solutions”. In: *AgriEngineering* 4.1 (2022), pp. 70–103. DOI: <https://doi.org/10.3390/agriengineering4010006>.
- [2] Paolo Bertellini et al. “Binary Classification of Agricultural Crops Using Sentinel Satellite Data and Machine Learning Techniques”. In: *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 859–864. DOI: 10.15439/2023F1703.
- [3] Nicolas R Dalezios et al. “Remotely sensed methodologies for crop water availability and requirements in precision farming of vulnerable agriculture”. In: *Water resources management* 33 (2019), pp. 1499–1519. DOI: <https://doi.org/10.1007/s11269-018-2161-8>.
- [4] Jonathan P Dash, Grant D Pearse, and Michael S Watt. “UAV multispectral imagery can complement satellite data for monitoring forest health”. In: *Remote Sensing* 10.8 (2018), p. 1216. DOI: 10.3390/rs10081216.
- [5] Patrick Ebel et al. “SEN12MS-CR-TS: A Remote Sensing Data Set for Multi-modal Multi-temporal Cloud Removal”. In: *IEEE Transactions on Geoscience and Remote Sensing* (2022). DOI: 10.1109/TGRS.2022.3146246. URL: <https://doi.org/10.1109/TGRS.2022.3146246>.
- [6] Patrick Ebel et al. “UnCRtainTS: Uncertainty quantification for cloud removal in optical satellite time series”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, pp. 2086–2096. DOI: 10.1109/CVPRW59228.2023.00202.
- [7] Kenji Enomoto et al. “Filmy cloud removal on satellite imagery with multispectral conditional generative adversarial nets”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 2017, pp. 48–56. DOI: 10.1109/CVPRW.2017.197.
- [8] Jichong Han et al. “Prediction of winter wheat yield based on multi-source data and machine learning in China”. In: *Remote Sensing* 12.2 (2020), p. 236. DOI: <https://doi.org/10.3390/rs12020236>.
- [9] Jay Ram Lamichhane et al. “Abiotic and biotic factors affecting crop seed germination and seedling emergence: a conceptual framework”. In: *Plant and soil* 432 (2018), pp. 1–28. DOI: <https://doi.org/10.1007/s11104-018-3780-9>.
- [10] Konstantinos G Liakos et al. “Machine learning in agriculture: A review”. In: *Sensors* 18.8 (2018), p. 2674. DOI: 10.3390/s18082674.
- [11] David B Lobell, Kenneth G Cassman, and Christopher B Field. “Crop yield gaps: their importance, magnitudes, and causes”. In: *Annual review of environment and resources* 34 (2009), pp. 179–204. DOI: <https://doi.org/10.1146/annurev.enviro.041008.093740>.
- [12] Andrea Meraner et al. “Cloud removal in Sentinel-2 imagery using a deep residual neural network and SAR-optical data fusion”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 166 (2020), pp. 333–346. DOI: <https://doi.org/10.1016/j.isprsjprs.2020.05.013>.
- [13] Michele Meroni et al. “Yield forecasting with machine learning and small data: What gains for grains?” In: *Agricultural and Forest Meteorology* 308 (2021), p. 108555. DOI: <https://doi.org/10.1016/j.agrformet.2021.108555>.
- [14] José M Peña-Barragán et al. “Object-based crop identification using multiple vegetation indices, textural features and crop phenology”. In: *Remote Sensing of Environment* 115.6 (2011), pp. 1301–1316. DOI: 10.1016/j.rse.2011.01.009.
- [15] Darius Phiri et al. “Sentinel-2 data for land cover/use mapping: A review”. In: *Remote Sensing* 12.14 (2020), p. 2291. DOI: <https://doi.org/10.3390/rs12142291>.
- [16] Francis J Pierce and Peter Nowak. “Aspects of precision agriculture”. In: *Advances in agronomy* 67 (1999), pp. 1–85. DOI: [https://doi.org/10.1016/S0065-2113\(08\)60513-1](https://doi.org/10.1016/S0065-2113(08)60513-1).
- [17] Yinghui Quan et al. “A novel feature extension method for the forest disaster monitoring using multispectral data”. In: *Remote Sensing* 12.14 (2020), p. 2261. DOI: <https://doi.org/10.3390/rs12142261>.
- [18] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: (May 2015). DOI: 10.48550/ARXIV.1505.04597. arXiv: 1505.04597 [cs.CV].
- [19] Marc Rußwurm and Marco Körner. “Self-attention for raw optical satellite time series classification”. In: *ISPRS journal of photogrammetry and remote sensing* 169 (2020), pp. 421–435. DOI: <https://doi.org/10.1016/j.isprsjprs.2020.06.006>.
- [20] Vivien Sainte Fare Garnot and Loic Landrieu. “Lightweight Temporal Self-Attention for Classifying Satellite Images Time Series”. In: *arXiv preprint arXiv:2007.00586* (2020). DOI: <https://doi.org/10.48550/arXiv.2007.00586>.
- [21] Mark Sandler et al. “Mobilenetv2: Inverted residuals and linear bottlenecks”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 4510–4520. DOI: 10.1109/CVPR.2018.00474.
- [22] Vishnu Sarukkai et al. “Cloud removal from satellite images using spatiotemporal generative networks”. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2020, pp. 1796–1805. DOI: 10.1109/WACV45572.2020.9093564.
- [23] Mohsen Shahhosseini et al. “Coupling machine learning and crop modeling improves crop yield prediction in the US Corn Belt”. In: *Scientific reports* 11.1 (2021), p. 1606. DOI: <https://doi.org/10.1038/s41598-020-80820-1>.
- [24] Ramon Torres et al. “GMES Sentinel-1 mission”. In: *Remote sensing of environment* 120 (2012), pp. 9–24. DOI: <https://doi.org/10.1016/j.rse.2011.05.028>.

- [25] Amanda Veloso et al. “Understanding the temporal behavior of crops using Sentinel-1 and Sentinel-2-like data for agricultural applications”. In: *Remote sensing of environment* 199 (2017), pp. 415–426. DOI: <https://doi.org/10.1016/j.rse.2017.07.015>.
- [26] Jonathan Weekley, Joseph Gabbard, and Jerzy Nowak. “Micro-level management of agricultural inputs: emerging approaches”. In: *Agronomy* 2.4 (2012), pp. 321–357. DOI: <https://doi.org/10.3390/agronomy2040321>.
- [27] Fang Xu et al. “GLF-CR: SAR-enhanced cloud removal with global–local fusion”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 192 (2022), pp. 268–278. DOI: <https://doi.org/10.1016/j.isprsjprs.2022.08.002>.
- [28] Qiangqiang Yuan et al. “Deep learning in environmental remote sensing: Achievements and challenges”. In: *Remote Sensing of Environment* 241 (2020), p. 111716. DOI: <https://doi.org/10.1016/j.rse.2020.111716>.
- [29] Chongyuan Zhang, Afef Marzougui, and Sindhuja Sankaran. “High-resolution satellite imagery applications in crop phenotyping: An overview”. In: *Computers and Electronics in Agriculture* 175 (2020), p. 105584. DOI: <https://doi.org/10.1016/j.compag.2020.105584>.
- [30] Zhe Zhu, Shixiong Wang, and Curtis E Woodcock. “Improvement and expansion of the Fmask algorithm: Cloud, cloud shadow, and snow detection for Landsats 4–7, 8, and Sentinel 2 images”. In: *Remote sensing of Environment* 159 (2015), pp. 269–277. DOI: <https://doi.org/10.1016/j.rse.2014.12.014>.

Organizational Capabilities for Business-IT Integration in Digital Enterprises

Constanze Riedinger
0009-0003-0226-4114
kips, Konstanz University of Applied Sciences,
78467 Konstanz, Germany
Email: constanze.riedinger@htwg-konstanz.de

Maike Netscher
0000-0003-0075-3339
Neu Ulm University of Applied Sciences,
89231 Neu-Ulm, Germany
Email: maike.netscher@hnu.de

Stephan Zimmermann
0009-0000-1977-6659
Technical University of Applied Sciences Augsburg,
86161 Augsburg, Germany
Email: stephan.zimmermann@tha.de

Abstract—The digital transformation urges organizations to become digital enterprises. Digital enterprises require the integration of business and IT to efficiently leverage digital technologies. However, there is a lack of a framework that guides organizations on what organizational capabilities are required to achieve business-IT integration. The goal of this paper is to identify these capabilities. From a single case study, we derived twelve organizational capabilities that a digital enterprise, driven by technology, should design in terms of its people, organizational structure, and tasks to establish business-IT integration. Thus, this paper provides guidance for organizations to approach business-IT integration as a foundation for their path into a digital enterprise.

Index Terms—Digital enterprise, organizational capabilities, case study, business-IT integration, Leavitt’s diamond model.

I. INTRODUCTION

IN 2017 digital transformation was the key topic of the World Economic Forum in Davos with leading business and political participants [1]. This transformation process, in which digital technologies are causing disruption, is triggering strategic responses from companies that try to adapt their value creation and manage changes in their structures to achieve the best outcomes in this process [2]. While digital technology has changed rapidly in the meantime, with upcoming topics like Artificial Intelligence, Low-Code Development, Cloud Computing, Internet of Things (IOT), etc. [3], companies are still struggling to achieve this transformation on a broad organizational level [4]. The results of a recent survey from 2024 indicate that the topics of organizational change, agile organization, and business-IT collaboration remain at the top of CIOs’ agendas [5]. Digital maturity models identify the key issue of digital enterprises [6] as a part of digital transformation that should be responsible for parallelizing and integrating organizational with technological change, or reducing the gap in the rate of change between the two. Furthermore, regarding the outcome of digital transformation, new technologies themselves do not bring competitive advan-

tages, however, superior management of technology does [7]. In the vision to become an integrated technology organization, companies must, therefore, identify measures to manage their digital transformation. Current research is focused on determining e.g. the maturity of IT management [8], business process management [9], or knowledge management [10] in digital enterprises. To foster these measures, however, companies still struggle with the adoption of their organization to the digital technology changes and in operationalizing these approaches on a broad level [11]. According to this problem, we aim to bring in a more strategic perspective on these measures.

Literature shows that one way to structure organizational challenges in a technology environment and in a way to receive requirements for their operational resolution is the usage of capabilities [12]. The organizational capabilities of a company are crucial to its competitiveness and growth [11]. They enable the formulation and implementation of strategies that build on the organization’s strengths and resources, and the successful deployment of competitive strategies that enable it to survive and grow in value over the long term [11]. Our goal of the paper is to position this idea of a capability-based approach and derive organizational capabilities for a business-IT integration based on a case study approach, leading to our research question: *What organizational capabilities should digital enterprises design to establish business-IT integration?*

As a result, we come up with organizational capabilities in the fields of *structure, technology, people, and task* following Leavitt’s Diamond model [13]. These capabilities can help managers in an enterprise driven by technology to define what they need to do to develop people as cornerstones and build up collaborative tasks and shared-responsibility structures. This fosters business-IT integration in their digital enterprise.

Based on the background of digital enterprises, business-IT integration, and organizational capabilities in Chapter II, and our case study research method in Chapter III, our findings provide an overview of organizational capabilities

for business-IT integration in Chapter IV. We conclude with a discussion and implications for theory and practice.

II. THEORETICAL BACKGROUND

A. *Digital enterprises*

A digital enterprise describes an organization that recognizes digital transformation as a fundamental element of its strategy, culture, and operations to succeed in the digital age [14]. It employs digital technologies to achieve a sustainable, technology-based competitive advantage [15]. However, becoming a digital enterprise involves not just adopting technology but fundamentally changing how the organization operates [7]. Consequently, it is also characterized by a quick adaptability to changing market conditions and customer needs, new ways of collaboration among employees and with partners as well as management practices that emphasize self-organization and proactive leadership [14]. Now, that digital transformation and the rapid development of innovative technologies are forcing companies to incorporate digitalization into their value-creation process, the transformation to a digital enterprise is imperative to remain competitive [16].

In this context, digital maturity refers to the level of readiness for digital transformation. Several digital maturity models are established for enterprises to assess and benchmark their digital readiness [15]. They provide a structured framework for evaluating the organization's current state based on specific criteria, identifying areas for improvement, and defining a roadmap for advancing digital abilities [7]. Besides technologies, infrastructure, and operations, strategy, organizational structure, and people in the company are key components when transforming into a digital enterprise [16]. However, for a more profound comprehension and operationalization, the fusion of business and IT is crucial for digital enterprises: the development of a digital business strategy necessitates a merger of business and IT strategies [17]. An efficient use of digital technologies and the implementation of IT innovations require the integration of business and IT [18].

B. *Business-IT Integration*

The interplay between business and IT has been extensively studied for several decades [18], [19]. An established model to explain this interplay is the Strategic Alignment Model (SAM) by Henderson and Venkatraman [20], serving also as a basis for future research in the field of business-IT alignment (BITA) [17]. It describes alignment as a multivariate relationship based on the fit between strategy and operative domains, as well as the functional integration between business and IT domains [20]. Alignment, however, is not a final state but rather described as a dynamic and evolutionary process [21]. Therefore, organizations need to consider criteria such as governance, partnership, scope and architecture, skills, value measurement as well as communication to improve BITA [21]. Enterprises strive for consensus between business and IT, as successful alignment has been proven to lead to better performance [19].

The imperative for organizations to turn into a digital enterprise is leading academics to rethink the basic idea of aligning business and IT as separate departments towards a fusion [17] or integration [18]. At the very least, as shown by alignment literature and regardless of the structure, with or without a separate IT function, organizations require the merging of business and IT strategy and convergence at the operational level [17], leading to the integration of business (as business models, products, capabilities, and processes) and IT (as technology). In the context of a digital enterprise, this integration implies considering different perspectives such as "technologies, employees, management and social elements" [18] according to the socio-technical understanding.

C. *Organizational Capability Framework*

In order to achieve the integration of business and IT in digital enterprises, it is essential to develop the necessary organizational capabilities [12]. Following [22], these organizational capabilities are the ability to manage the organizational structure, the technologies, and the human capabilities. This encompasses, for instance, the information technologies employed in the business units and the skills of the people working in these units in line with the strategic goals. Taking these organizational aspects into account is crucial for the integration of business and IT [22]. Moreover, organizational capabilities are defined as the ability of an organization to perform coordinated tasks with its resources to achieve strategic goals, manage activities more effectively, and keep pace with digital transformation [23], [24]. The summary of these required organizational capabilities is referred to as a capability framework [25]. Capability frameworks are conceptual maps for systematically capturing, organizing, and developing the required skills [26]. In more specific terms, capability frameworks delineate the competencies, knowledge, tools, processes, and behaviors, thereby elucidating the fundamental elements of value creation [27]. The objective of them is to achieve a higher level of digital maturity [12], [28].

Reference [22] proposed one of the earliest organizational capabilities frameworks regarding the broader topic of digital transformation. The most crucial capabilities for a digital change of enterprises were identified as abilities of coordination and integration through efficient collaboration and information sharing across different departments such as IT and business units [22]. In addition, the ability to scale, the ability of managers to engage in continuous education, and the efficient management of new technologies are important capabilities [22]. Furthermore, [23] developed themes of organizational capabilities that arise in the context of digital transformation. In terms of the capabilities required for business-IT integration, three important capabilities are mentioned: innovative thinking, organizational design, and digital transformation leadership [23].

Further, uncoupled research presents several factors that influence the digital maturity of an organization regarding the integration of business and IT: Reference [29] elaborate patterns that describe digitally mature organizations including the investment in new technologies as well as the education

of employees. Reference [30] studies strategic factors that influence digital maturity and highlight the shared vision of top management concerning digitalization and transformation as important factors. Reference [31] outlines that especially the characteristics and competencies of digital leadership influence digital maturity. Reference [32] describes the importance of governance in IT as essential. Key practices in this study include integrated IT decision leadership at the executive and board decision levels, and the establishment of digital leaders who act for both business and IT. In addition, reference [33] suggests that architectural thinking, including collaborative management of enterprise architecture and value-based service, leads to improved maturity. Reference [34] identifies also digital leadership, collaboration culture, innovation culture, and agile processes as organizational capabilities for digital maturity. In particular, the prioritization of resources aligned with operational needs and strategic goals, along with the implementation of flexible and responsive decision-making processes, can facilitate the achievement of agile processes [35], [11]. Another relevant factor to achieve collaboration between business and IT counterparts is their communication based on a common language [36]. Despite the existence of these valuable studies, there is a lack of capability frameworks that explicitly focus on the organizational capabilities required to influence an organization's digital maturity through the integration of business and IT, and that are comprehensively derived from practice.

III. METHODOLOGY

Organizational capabilities are a practical construct established in organizations and influenced by the organizational context [22], [23]. Therefore, we answer our research question by applying a case study approach. It ensures the practical perspective, the in-depth analysis as well as the exploration of contextual factors that are required for the analysis in the context of the integration of business and IT in organizations [21]. Furthermore, the case study approach is recommended to research a contemporary phenomenon [37] such as the digital enterprise and has been applied in previous IS research on capabilities [27]. By collecting data from various sources during the case study, we enrich our insights and the validity of our findings [38].

A. Case Study Design

We follow case study research guidelines to design our case study and to select our case [37], [39]. We chose a single-case design with a common case [37], representative for most global companies: Our case was a major engineering enterprise with a size of about 6000 employees and the existence of all standard business areas including several divisions based on different products. We could accompany the case for one year and observe their path to become leader in innovation in their field, pursuing smart, advanced and cyber secure solutions and embracing new technology trends such as artificial intelligence, big data analysis and IOT to achieve it. In order to study the business-IT integration in depth and to enhance the insights, we focused on multiple units of analysis as

proposed by [37] in a single case embedded design. The organization's operating model indicated a distinction by divisions and departments leading to twelve units of analysis. Some of them had sub-departments which we included in the data collection, but then grouped in the analysis based on the higher hierarchical level.

B. Data Collection

Different sources of evidence, mainly documents, interviews, and focus groups served the data collection [37]: We acquired various strategy documents and operating models as a basis for the analysis. Based on the strategic vision of the IT department to enable the company's strategy through a strong business focus and the provision of optimal digital solutions, we conducted a workshop with managers from the IT department. In this focus group workshop, the target was to analyze the current strategic vision regarding business-IT integration and to discuss the readiness to become a business partner providing real business value. Following the identified units of analysis derived from the organizational structure, we also conducted 38 interviews with responsible managers of different divisions, departments, and sub-departments. We discussed how IT can support or enable their business capabilities to achieve their strategic goals and how the integration of IT into their business should look like. Furthermore, we collected insights in documents from the enterprise architecture (EA) such as EA principles, business capabilities, applications, and IT processes, and interviewed the responsible enterprise architect to gain insights into the interrelations of business needs and IT within the organization. The following list presents a detailed overview of the various data sources:

- Overall organizational strategy and operating model
- 12 strategies for each unit of analysis
- IT department strategy, its operating model, and IT processes
- 38 interviews with business managers of divisions and departments following the organizational structure
- Workshop with IT managers (CIO, IT Service Management, IT Infrastructure and Applications, IT Business Partner Management)
- Enterprise architecture principles and enterprise architecture assets including business capabilities as well as processes
- Interview with the enterprise architect

C. Data Analysis and Synthesis

We analyzed the documents acquired in the case study as a first step of data analysis. Three researchers participated in this process. We conducted qualitative coding [40] based on our research question and extracted tasks, success factors, and goals emphasized in strategy documents and the operating models. We further analyzed the notes taken during interviews and the IT workshop. Thereby, we took into account the collaboration prevailing in the organization and the pursued new ways of collaboration that are characteristic for a digital enterprise [14]. We derived tasks and requirements to

achieve new ways of collaboration and support the digital enterprise through business-IT integration. We then examined IT processes as well as artifacts of enterprise architecture management (EAM) to understand current IT solution delivery and the interrelations of business and IT. Based on the qualitative analysis, we derived organizational capabilities required for business-IT integration.

We then synthesized the derived organizational capabilities based on an existing framework. As shown in earlier studies [36], [41], this alignment ensures not only the link to an existing framework but also a more comprehensive approach [42]. Therefore, we chose Leavitt's Diamond model [13] which is a change management framework that serves the analysis of causal relationships within an organization and the development of adaptation strategies. It has been applied in earlier studies to analyze the various changes through digital transformation and required adaptations of organizations [43], [44]. Thus, it also fits as an underlying framework for our research as we aim to study digital enterprises impacted by digital transformation and their required capabilities to integrate business and IT in this changing context.

Leavitt's Diamond model considers four key dimensions in organizations - people, structure, task, and technology - that influence each other [13]. The theory is that a change in one of these dimensions will inevitably lead to a change in all the others. The balance of all dimensions is crucial for change to work and is the basis for an organization's ability to respond to change and achieve its goals. [45] Accordingly, the disruption through innovative technologies in digital enterprises requires immediate change and adaption in the other three dimensions. Based on this assumption, [44] adjusted Leavitt's Diamond model with technology in the center having a strong impact on the other dimensions but also highlighted the strong interrelation of all dimensions. This leads to a new conceptual model, in which technology is the key driver for organizational transformation. We followed those assumptions of technology as key driver and applied selective coding [40] to assign the identified organizational capabilities to the three dimensions [13]: *people*, *structure*, and *task*. *People* are all the individuals within the organization, including their skills, knowledge, attitudes, and interactions. *Structure* describes the organizational design and hierarchy that defines how tasks are divided and coordinated within the organization. *Tasks* include all specific activities and responsibilities performed by individuals or groups to achieve organizational objectives [13].

IV. FINDINGS

This section presents the findings of the case study. Fig. 1 displays the organizational capabilities clustered by the Leavitt's diamond dimensions *people*, *structure*, and *task*. As outlined in the methodology section, *technology* is the key driver and, therefore, placed in the center.

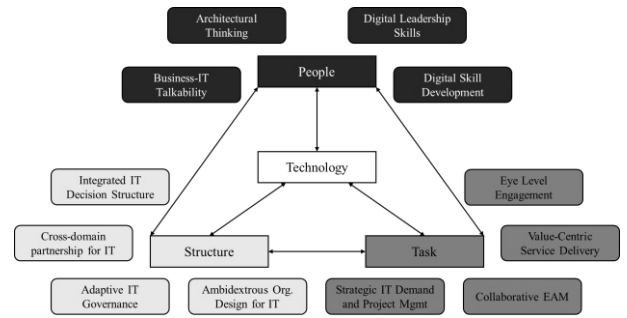


Fig. 1. Organizational Capabilities clustered by Leavitt's Diamond Dimensions

Furthermore, the organizational capabilities are described based on the analysis of the collected data in Table 1. They are clustered by the Leavitt's Diamond Dimensions and the description outlines what an organization must do and what can be achieved by designing each capability. Those descriptions of the organizational capabilities summarize the qualitative codes of our data analysis.

TABLE 1. ORGANIZATIONAL CAPABILITIES AND DESCRIPTIONS CLUSTERED BY LEAVITT'S DIAMOND DIMENSIONS

	Organizational Capability	The organization should be able to...
People	Business-IT Talkability	...empower individuals across departments to communicate in one language about business and IT, facilitating effective collaboration and understanding to drive value through IT in digital enterprises.
	Architectural Thinking	...foster individuals with the mindset and skills to align technology with business components, link technological interrelations for strategic value creation in digital enterprises, and understand the benefits and goals of structuring and managing an EA.
	Digital Leadership Skills	... empower individuals to embrace change, leverage technology, and collaborate innovatively to drive digital transformation and achieve strategic objectives in the organization.
	Digital Skill Development	... empower individuals with the technical and business acumen necessary to bridge the gap between IT abilities and organizational objectives, driving collaboration and innovation.
Task	Strategic IT Demand and Project Management	... prioritize and align resources with operational needs and strategic objectives, to optimize resource allocation and thereby facilitate seamless integration between business and IT.
	Collaborative Enterprise Architecture Management	... facilitate coordinated efforts among stakeholders to develop and maintain a comprehensive architectural framework that aligns business strategies with IT abilities and ensures a jointly developed target landscape for the future of the digital enterprise.
	Value Centric Service Delivery	... prioritize customer-centricity and strengthen the service focus by and in delivering IT solutions with a culture of responsiveness, efficiency, and value creation. This may also include empowering citizen development enabling various stakeholders for (self-)service delivery.
	Eye level Engagement	... facilitate peer-to-peer communication on eye level and engage in one-to-one dialogues between stakeholders. This should consolidate issues, reflect IT needs and emphasize the value of IT in meeting business objectives to foster collaborative

		communities, overcome hierarchies, drive innovation, and ensure continuous alignment between business goals and IT initiatives.
Structure	Integrated IT Decision Structure	...align IT decision-making with broader organizational goals and strategies, ensuring IT representation at the board level and fostering collaboration between IT leaders and business executives to drive strategic alignment and maximize the effectiveness of digital initiatives within the enterprise.
	Cross-domain Partnership for IT	...establish visionary leaders who serve as advocates for both business and IT, integrating user and strategic demands across the organization to ensure organizational-wide alignment and future readiness.
	Adaptive IT Governance	... implement flexible and responsive decision-making structures and IT task responsibility allocation that adjust to changing internal and external dynamics, ensuring effective management and alignment of resources with strategic objectives.
	Ambidextrous Organizational Design for IT	...maintain and balance both stable, exploitative and dynamic, explorative organizational structures and processes for IT in response to changing market conditions, technological advancements, and customer needs, ensuring agility and resilience in the face of uncertainty. This may also include introducing communities for knowledge exchange, agile work structures, or structured dialogue formats.

The analysis revealed four capabilities for each of the dimensions of *people*, *structure*, and *task* that are crucial to establish business-IT integration. The *people* dimension emphasizes the importance of communication between business and IT in a **unified language**, **architectural thinking** in a collaborative culture, **digital leadership**, and the continuous development of **digital skills**. The *task* dimension encompasses the capability of **strategic IT demand and project management** to efficiently prioritize and allocate resources for IT. The **collaborative management of the enterprise architecture** enables the integration of business strategies and IT abilities in the IT landscape and its evolution. The **delivery of value-based IT services** focuses on customers and users by establishing a culture of collaboration on IT solutions and empowering individuals. Moreover, the ability to foster **engagement at an equal level** serves to facilitate open communication and collaborative communities. In the *structure* dimension, we identified the capability for **integrated IT decision-making** to elevate IT issues to board level. Furthermore, **cross-domain partnerships** aim for visionary digital leaders to advocate digital topics. An **adaptive IT governance** enables the implementation of flexible and responsive decision-making structures and deals with the allocation of IT task responsibilities between business and IT roles. The ability to establish an **ambidextrous organizational design** enables dynamic response to changes ensuring agility and resilience. These organizational capabilities are crucial for effective collaboration, embracing technological change, integrating IT with business goals, and driving digital transformation.

V. DISCUSSION

The findings of the case study answer the research question: *What organizational capabilities should digital enterprises design to establish business-IT integration?* In the following, we discuss our findings on this question regarding Leavitt's Diamond Dimensions in terms of practical implications and existing research.

A. Technology as Key Driver in Digital Enterprises

Digital disruption leads to an imbalance of the Leavitt's Diamond [13] and requires organizations to design organizational capabilities in the other three dimensions in order to achieve business-IT integration and become a digital enterprise. We have identified these organizational capabilities, with the technology dimension considered as an influencing factor that needs to be balanced and supported by the people, structure, and task dimensions. The identified capabilities from each dimension result in an overall construct as a framework that interacts to influence the maturity of business-IT integration. The maturity of the capabilities required in the three dimensions depends on the intensity and type of technology used as well as the strategic goals and size of the organization. For instance, it can be assumed that smaller digital enterprises whose value creation is based on the use of IT require a more intensive development of the capabilities in the dimensions than large enterprises in other sectors. In accordance with [44], our framework demonstrates that technology is a key driver of enterprise changes. However, this does not negate the influence of the development of capabilities in other areas on the selection and management of technology.

B. People as a Cornerstone for Business-IT Integration

The people in an organization are an important cornerstone for the success of change in the company [11]. A digital enterprise, therefore, needs digital leaders who drive the digital transformation and lead the way [31]. The case study shows the need for **digital leadership** in various interviews with business executives. Those who embraced change and leveraged technology could drive innovation and worshiped a close relationship to IT colleagues to shape change together. However, in one interview, the IT colleague emphasized that *the collaboration succeeded not only because the business head of customer service led the way as a digital leader, but also because the communication between business and IT worked and a common language was found between business and IT*. This created a common understanding of how the technologies can be used in the best possible way. This **business-IT talkability** and the use of a common language is also reflected in earlier studies on the alignment and successful communication [36] and communication patterns in project teams [46]. **Digital skills** in both business and IT are important for the digital enterprise. In the case study, a strategic pillar of the IT department was to enable growth by actively managing the architecture. This requires an organizational capability of **architectural thinking** where people understand the links between technology and business goals and get the big picture

of the enterprise architecture. However, the lack of architectural thinking within the organization, especially among the business employees, became apparent by interviewing the enterprise architect. He highlighted, as mentioned by prior research [33] that *especially EAM communication is required to create awareness for the skills needed to collaboratively design the companies' architecture.*

C. Tasks based on Collaboration for Business-IT Integration

The task to perform *collaborative EAM* is an important organizational capability for the digital enterprise [33]. Studies show that the potential of EAM can only be exploited if it is driven by shared responsibility [33]. The aim of the case study was to bring stakeholders together through EAM and jointly develop the future landscape regarding scarce resources and the most value-adding IT support. Value centrality was also highlighted in the case as main pillar of the companies' strategy. This also influenced the collaboration between business and IT regarding the *delivery of services*. To implement digitalization projects, the corporate culture is crucial for success [34], [47]. Therefore, the stakeholders especially from IT department were encouraged to put the customer at the center and work with them to drive value creation forward. However, for business-IT integration joint teams should possess the ability to deliver IT solutions with a clear service focus. And also, business stakeholders could be empowered to contribute to IT service delivery as citizen developers [48].

From the focus on digitalization projects to become a leader in innovation arose the necessity for further organizational capabilities promoting business-IT integration. Therefore, those tasks were optimized during the one-year support. First, the ability to transparently define a demand and project portfolio process required to have a long-term focus on IT projects. Such a *strategic IT demand and project management* ensures that operational needs and strategic objectives are aligned and prioritized [18]. In the case, this involved teams consisting of business and IT working at eye-level together to decide on demands and digitalization projects and allocate resources at an early stage. Second, the ability to communicate openly and work together was essential to become a digital enterprise. A regular discussion between business and IT as a *collaborative engagement* supports business-IT integration [36] and, therefore, marks an organizational capability.

D. Structures support shared-responsibility for Business-IT Integration

Organizational structures are required to perform those tasks and drive the development of becoming a digital enterprise [23]. Digital enterprises need structures and processes that can be quickly adapted to environmental influences [35] - i.e. the organizational capability for *ambidextrous organizational design with regard to IT*. Following the operating model in the case study, the company initiated various restructurings to reflect this flexibility, particularly regarding

the innovative product divisions. To thereby implement and strengthen the integration of business and IT, a structured dialogue between diverse stakeholders was established and communities for knowledge sharing were considered. The flexibility in structures and processes must also be ensured in decision-making [35]. Managers in the case study discussed, regarding IT developed in business departments, that *adaptive IT governance would reflect a modern way to address a flexible allocation of responsibilities*. This type of responsive decision-making and allocation of responsibilities supports the interaction between different stakeholders and the integration of business and IT [18].

Besides responsive decision-making processes, a common notion for a digital enterprise is that *IT decision structures* must be integrated into the company structures [11]. The structures should reflect IT decision-making at board level, which is why it is recommended to incorporate a chief information officer or chief digital officer as an advocate for IT and digital initiatives within the board. In the case, it became apparent that hierarchical structures did not allow this direct representation at board level. Therefore, we identified that to become a digital enterprise further restructuring of the operating model would be recommended. However, not only alignment on the executive level but also the operational level is required for integration [18]. Therefore, *digital leaders* who serve as advocates for both business and IT should be established throughout the organization and drive transformation. Such a *cross-domain partnership for IT* may integrate user and strategic demands across the organization and foster business-IT integration [18].

E. Balancing the dimensions of the digital enterprise

The Leavitt's Diamond model states that there is an interaction between the four dimensions [13]. Not only does technology as a key driver influence the other dimensions, but they also influence each other [13]. Those interrelations between the organizational capabilities, clustered in different dimensions of the Leavitt's Diamond, became apparent during the case study: Some tasks could only be carried out with corresponding structures in place and the right skills available among the employees. For *engagement on eye level*, they established a structured dialogue format and trained responsible IT managers with *business talk* skills to jointly emphasize and shape the added value of IT for the business. To perform *collaborative EAM*, the company still lacked employees who incorporate *architectural thinking* and corresponding responsibilities that are clearly defined by an *adaptive IT governance*.

The clustering of the identified organizational capabilities into the dimensions of the Leavitt's diamond represents a structuring for our capability framework and stresses the interplay of the capabilities, however, it does not imply any delimitation of content. Following prior research, digital leadership marks a crucial capability of a digital enterprise [34]. However, to establish business-IT integration within a digital enterprise, we identified that *digital leadership* has two dimensions to consider: it necessitates leadership skills in business and IT to drive innovation, but it also requires structures

that represent this leadership and enable transformation and strategic alignment.

VI. CONCLUSION

Digital transformation urges companies to transform into digital enterprises to achieve maximal value from technology and stay competitive. The integration of business and IT is essential for a digital enterprise to effectively and efficiently use digital technologies and implement innovations. However, there is no framework guiding organizations on what capabilities are required to achieve business-IT integration. Therefore, this paper identifies these organizational capabilities that digital enterprises need to design to establish business-IT integration. To reach this goal, we use a single case study of a representative company that we have followed for a year as it develops into an innovation leader in its field. We identified twelve organizational capabilities required to establish business-IT integration. Those capabilities relate to the dimensions of the Leavitt’s diamond of *people*, *task*, and *structure* whereas *technology* is the key driver for change. To become a digital enterprise and thereby achieve business-IT integration, organizations need to design those capabilities, manage their interrelationships, and balance the dimensions of the Leavitt’s Diamond.

This paper contributes to research by enhancing research on organizational capabilities linked to the field of collaboration between business and IT. By applying the Leavitt’s Diamond, it also advances research on organizational change management and the interdependencies of the dimensions when it comes to business-IT integration. For practitioners, the organizational capabilities offer guidance on what their company must design to achieve business-IT integration and thereby lay the foundation to become a digital enterprise.

Nevertheless, the study has limitations that provide avenues for future research: first, the design of organizational capabilities is a first step for business-IT integration that still lacks operationalization. To address this limitation, researchers could study with focus groups or in further case studies how the identified capabilities may be achieved and operationalized. Thereby, the prevailing corporate culture of an enterprise which was noted in several capabilities should be considered as an important factor for the operationalization. Furthermore, they could analyze which measures or key performance indicators are required to assess these capabilities and design them in new organizational settings to determine long-term effects on business-IT integration. Second, the required degree of these capabilities is contingent upon the intensity of the technological influence. Future research could consider the type of enterprise and the required use of technologies in line with the capability framework.

REFERENCES

- [1] World Economic Forum, ‘Annual Meeting 2017 Responsive and Responsible Leadership’, Davos-Klosters, Switzerland, 2017. Accessed: May 26, 2024. [Online]. Available: https://www3.weforum.org/docs/WEF_AM17_Report.pdf
- [2] G. Vial, ‘Understanding digital transformation: A review and a research agenda’, *The Journal of Strategic Information Systems*, vol. 28, no. 2, pp. 118–144, Jun. 2019, <http://dx.doi.org/10.1016/j.jsis.2019.01.003>.
- [3] Gartner Inc., ‘Top 10 Strategic Technology Trends for 2024’. Accessed: May 23, 2024. [Online]. Available: <https://www.gartner.com/en/articles/gartner-top-10-strategic-technology-trends-for-2024>
- [4] E. Obu-Cann, G. Fletcher, M. Griffiths, M. Kutar, and S. Krishnan-Harihara, ‘Digital Maturity and SMEs: evaluating the application of a digital maturity assessment tool’, 2023.
- [5] Gartner Inc., ‘Technology Priorities CIOs Must Address in 2024’. Accessed: May 23, 2024. [Online]. Available: <https://www.gartner.com/en/confirmation/information-technology/research/cio-agenda-thank-you-it>
- [6] G. C. Kane, D. Palmer, N. Phillips, D. Kiron, and N. Buckley, ‘Achieving Digital Maturity’, *MIT Sloan Management Review and Deloitte University Press*, 2017.
- [7] D. L. Soule, A. Puram, G. F. Westerman, and D. Bonnet, ‘Becoming a Digital Organization: The Journey to Digital Dexterity’, *SSRN Journal*, 2015, <http://dx.doi.org/10.2139/ssrn.2697688>.
- [8] J. Becker, R. Knackstedt, and J. Poepelbuss, ‘Developing Maturity Models for IT Management’, *Business & Information Systems Engineering*, vol. 1, pp. 213–222, Jun. 2009, <http://dx.doi.org/10.1007/s12599-009-0044-5>.
- [9] M. Rohloff, *Case Study and Maturity Model for Business Process Management Implementation*, vol. 5701. 2009, p. 142. http://dx.doi.org/10.1007/978-3-642-03848-8_10.
- [10] U. Kulkarni and R. Freeze, *Development and Validation of a Knowledge Management Capability Assessment Model*. 2004, p. 670.
- [11] J. B. Akaegbu and A. A. Usoro, ‘The place of organisational capabilities in strategy formulation and implementation: an exploratory analysis’, *Glo J. Soc. Sci.*, vol. 16, no. 1, p. 39, Jan. 2018, <http://dx.doi.org/10.4314/gjss.v16i1.5>.
- [12] J. Van Riel, G. Poels, and S. Viaene, ‘Exploring Capability Mapping as a Tool for Digital Transformation: Insights from a Case Study’, in *Enterprise, Business-Process and Information Systems Modeling*, vol. 479, H. Van Der Aa, D. Bork, H. A. Proper, and R. Schmidt, Eds., in *Lecture Notes in Business Information Processing*, vol. 479., Cham: Springer Nature Switzerland, 2023, pp. 243–255. http://dx.doi.org/10.1007/978-3-031-34241-7_17.
- [13] H. J. Leavitt, ‘Applied organizational change in industry: Structural, technological and humanistic approaches’, *Handbook of organizations*, 1965.
- [14] C. C. Snow, Ø. D. Fjeldstad, and A. M. Langer, ‘Designing the digital organization’, *J Org Design*, vol. 6, no. 1, 2017, <http://dx.doi.org/10.1186/s41469-017-0017-y>.
- [15] L. Kontić and Đ. Vidicki, ‘Strategy for digital organization: Testing a measurement tool for digital transformation’, *Strategic Management*, vol. 23, no. 2, pp. 29–35, 2018, <http://dx.doi.org/10.5937/StraMan1801029K>.
- [16] W. Kersten, T. Blecker, and C. M. Ringle, ‘Adapting to the Future How Digitalization Shapes Sustainable Logistics and Resilient Supply Chain Management’, in *Proceedings of the Hamburg International Conference of Logistics (HICL)*, Berlin: epubli, 2021.
- [17] C. Kahre, D. Hoffmann, and F. Ahlemann, ‘Beyond Business-IT Alignment - Digital Business Strategies as a Paradigmatic Shift’, in *Proceedings HICSS 2017*, 2017.
- [18] C. Grawe, B. Kruse, and U. Bretschneider, ‘Business/IT Integration: Challenging the Boundaries of Alignment’, in *AMCIS 2020 Proceedings*, 2020, p. 20.
- [19] Y. E. Chan and B. H. Reich, ‘IT alignment: what have we learned?’, *J Inf Technol*, vol. 22, no. 4, pp. 297–315, 2007.
- [20] J. C. Henderson and H. Venkatraman, ‘Strategic alignment: Leveraging information technology for transforming organizations’, *IBM Systems Journal*, vol. 32, no. 1, pp. 472–484, 1993, <http://dx.doi.org/10.1147/sj.382.0472>.
- [21] J. Luftman and R. Kempaiah, ‘An Update on Business-IT Alignment: “A Line” Has Been Drawn’, *MIS Quarterly Executive*, vol. 6, no. 3, 2007.
- [22] A. D. Chandler and T. Hikino, *Scale and Scope: The Dynamics of Industrial Capitalism*. Harvard University Press, 1994.
- [23] J. Konopik, C. Jahn, T. Schuster, N. Hoßbach, and A. Pflaum, ‘Mastering the digital transformation through organizational capabilities: A conceptual framework’, *Digital Business*, vol. 2, no. 2, p. 100019, Jan. 2022, <http://dx.doi.org/10.1016/j.digbus.2021.100019>.

- [24] D. J. Collis, 'Research Note: How Valuable are Organizational Capabilities?', *Strat. Mgmt. J.*, vol. 15, no. S1, pp. 143–152, 1994, <http://dx.doi.org/10.1002/smj.4250150910>.
- [25] G. Bondel, A. Faber, and F. Matthes, 'Reporting from the Implementation of a Business Capability Map as Business-IT Alignment Tool', in *2018 IEEE 22nd International Enterprise Distributed Object Computing Workshop (EDOCW)*, Oct. 2018, pp. 125–134. <http://dx.doi.org/10.1109/EDOCW.2018.00027>.
- [26] Y. Malchenko, M. Gogua, K. Golovacheva, M. Smirnova, and O. Alkanova, 'A critical review of digital capability frameworks: a consumer perspective', *Digital Policy, Regulation and Governance*, vol. 22, no. 4, pp. 269–288, 2020, <http://dx.doi.org/10.1108/DPRG-02-2020-0028>.
- [27] R. Keller, O. Philipp, and R. Patrick, 'Pathways to Developing Digital Capabilities within Entrepreneurial Initiatives in Pre-Digital Organizations', *Business & Information Systems Engineering*, vol. 64, no. 1, pp. 33–46, Feb. 2022, <http://dx.doi.org/10.1007/s12599-021-00739-3>.
- [28] R. Cosic, G. Shanks, and S. B. Maynard, 'A business analytics capability framework', *Australasian Journal of Information Systems*, vol. 19, Sep. 2015, <http://dx.doi.org/10.3127/ajis.v19i0.1150>.
- [29] G. Gudergan, P. Mugge, A. Kwiatkowski, H. Abbu, T. L. Michaelis, and D. Krechting, 'Patterns of Digitization – What differentiates digitally mature organizations?', in *2019 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, Jun. 2019, pp. 1–8. <http://dx.doi.org/10.1109/ICE.2019.8792585>.
- [30] G. Salvioiti, A. Gaur, and F. Pennarola, 'STRATEGIC FACTORS ENABLING DIGITAL MATURITY: AN EXTENDED SURVEY', in *MCIS 2019 Proceedings*, 2019, p. 15.
- [31] H. Abbu, P. Mugge, G. Gudergan, and A. Kwiatkowski, 'DIGITAL LEADERSHIP - Character and Competency Differentiates Digitally Mature Organizations', in *2020 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, Jun. 2020, pp. 1–9. <http://dx.doi.org/10.1109/ICE/ITMC49519.2020.9198576>.
- [32] S. De Haes and W. Van Grembergen, 'An Exploratory Study into IT Governance Implementations and its Impact on Business/IT Alignment', *Information Systems Management*, vol. 26, no. 2, pp. 123–137, Apr. 2009, <http://dx.doi.org/10.1080/10580530902794786>.
- [33] B. Horlach, A. Drechsler, I. Schirmer, and P. Drews, *Everyone's Going to be an Architect: Design Principles for Architectural Thinking in Agile Organizations*. 2020. <http://dx.doi.org/10.24251/HICSS.2020.759>.
- [34] E. T. Devi, D. Wibisono, and N. B. Mulyono, 'Identifying critical capabilities for improving the maturity level of digital services creation process', *Journal of Industrial Engineering and Management*, vol. 15, no. 3, pp. 498–519, 2022, <http://dx.doi.org/10.3926/jiem.3818>.
- [35] G. M. Jonathan, L. Rusu, and W. Grembergen, 'Business-IT Alignment and Digital Transformation: Setting a Research Agenda', *Proceedings of the International Conference on Information Systems Development (ISD)*, Aug. 2021, [Online]. Available: <https://aisel.aisnet.org/isd2014/proceedings2021/currenttopics/7>
- [36] C. Riedinger, M. Huber, and N. Prinz, 'Factors for Effective Communication of IT Costs and IT Business Value', in *Proceedings of 18th Conference on Computer Science and Intelligence Systems (FedCSIS) 2023*, Sep. 2023, pp. 677–687. <http://dx.doi.org/10.15439/2023F7224>.
- [37] R. K. Yin, *Case study research: design and methods*, 5. edition. Los Angeles London New Delhi Singapore Washington, DC: SAGE, 2014.
- [38] U. Flick, *Triangulation*. Wiesbaden: VS Verlag für Sozialwissenschaften, 2011. <http://dx.doi.org/10.1007/978-3-531-92864-7>.
- [39] P. Baxter and S. Jack, 'Qualitative Case Study Methodology: Study Design and Implementation for Novice Researchers', *TQR*, Jan. 2015, <http://dx.doi.org/10.46743/2160-3715/2008.1573>.
- [40] J. M. Corbin and A. L. Strauss, *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*, 4th ed. Thousand Oaks, CA: Sage Publications, 2015.
- [41] N. Prinz, M. Huber, J. Leonhardt, and C. Riedinger, 'Unleash the Power of Citizen Development: Leveraging Organizational Capabilities for Successful Low-Code Development Platform Adoption', *HICSS 2024*, pp. 569–578, 2024.
- [42] M. Saunders, *Research Methods for Business Students*, 8th ed. Harlow, UK: Pearson, 2019.
- [43] D. Wigand, 'Building on Leavitt's Diamond Model of Organizations: The Organizational Interaction Diamond Model and the Impact of Information Technology on Structure, People, and Tasks', in *AMCIS 2007 Proceedings*,
- [44] J. Nograšek and M. Vintar, 'Technology as the Key Driver of Organizational Transformation in the eGovernment Period: Towards a New Formal Framework', in *Electronic Government*, vol. 6846, M. Janssen, H. J. Scholl, M. A. Wimmer, and Y. Tan, Eds., in Lecture Notes in Computer Science, vol. 6846, Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 453–464. http://dx.doi.org/10.1007/978-3-642-22878-0_38.
- [45] P. Kræmmergaard, B. C. Lyngø, and C. D. Schou, 'A Review of IT Service Management and Organisational Changes', 2008. [Online]. Available: <http://www.icesal.org/>
- [46] K. Muszyńska, 'Project Communication Management Patterns', in *Proceedings of Federated Conference on Computer Science and Information Systems (FedCSIS) 2016*, Oct. 2016, pp. 1179–1188. <http://dx.doi.org/10.15439/2016F235>.
- [47] C. Leyh, K. Köppel, S. Neuschl, and M. Pentrack, 'Critical Success Factors for Digitalization Projects', in *Proceedings of 16th Federated Conference on Computer Science and Information Systems (FedCSIS) 2021*, Sep. 2021, pp. 427–436. <http://dx.doi.org/10.15439/2021F122>.
- [48] S. Kanji, A. Katrodia, and F. Onyango Ogola, 'Effect of organizational capability on the performance of the financial services sector: A study of selected organizations in Kenya', *AJBER*, vol. 18, no. 2, pp. 51–71, 2023, <http://dx.doi.org/10.31920/1750-4562/2023/v18n2a3>.

Automated feedback generation in an intelligent tutoring system for counselor education

Eric Rudolph^{*¶}, Hanna Seer^{†¶}, Carina Mothes^{‡¶}, Jens Albrecht^{§¶}
 0009-0003-0615-4780^{*}, 0009-0006-1555-3603[†], 0009-0009-2930-0561[‡], 0000-0003-4070-1787[§]
 Nuremberg Institute of Technology Georg Simon Ohm[¶]

Abstract—This paper investigates the implementation of AI-driven feedback in an intelligent tutoring system (ITS) developed for training of counselors. By using LLMs, the study explores the automatic generation of feedback for communication-intensive tasks such as online counseling. The evaluation compares different feedback methods, including the sandwich, WWW and STATE methods, and assesses their emotional and objective impact. The results show that AI-generated feedback fulfills objective criteria better than emotional ones. Fine-tuning an open source LLM can improve both the emotional and objective quality of feedback. Furthermore, the study examines the acceptance of AI feedback among aspiring counselors, highlighting the influence of familiarity with AI on acceptance levels. Ethical considerations, including bias and hallucination, are addressed, with recommendations for risk mitigation through multi-feedback options and expert supervision. This research contributes to the understanding of the role of AI in improving digital counseling practices and highlights the need for continuous evaluation and ethical considerations.

I. INTRODUCTION

FEEDBACK can be a very effective learning tool [1]. However, it is a very communication-intensive and time consuming task for a teacher to provide feedback to every student. Especially in higher educational settings where classes sometimes have hundreds of students. Narciss et al. note that personalized tutoring feedback, especially using computer-based technologies, has significant potential for educational use [2]. This highlights the importance of leveraging technology to address the challenges of providing individualized feedback at scale.

To address this challenge, advancements in artificial intelligence (AI) offer promising solutions (i. e. [3], [4]). The rise of large language models (LLMs) marks a significant development in AI broadening its application across various domains, including mental health education. These technological advancements are particularly potent in enhancing communication skills through personalized AI-driven feedback mechanisms. One specific area where AI can play a transformative role is in the education of online counselors. Online counseling is a form of psycho-social support that is offered via the internet [5].

The shift from traditional face-to-face counseling to online modalities has been dramatically accelerated by the Covid-19 pandemic, which also increased the overall demand for psycho-social services [6]. This shift has underscored the need for effective training and feedback systems for practitioners in the digital counseling environment, where giving timely and

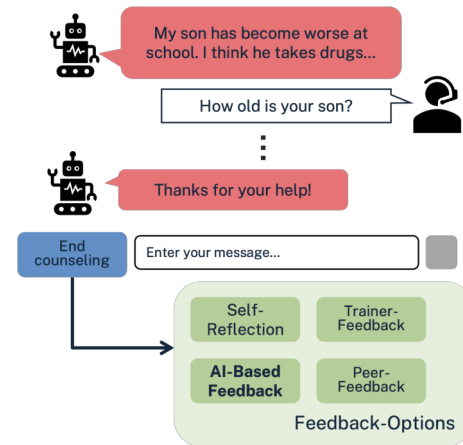


Fig. 1. Chat between a novice counselor and a chatbot pretending to be a client including feedback options after chat completion of our Intelligent Tutoring System

quality feedback can be both resource-intensive and challenging.

AI-based feedback in online counseling offers a promising support in counselor training, presenting unique opportunities to support and enhance the counseling process (i. e. in [7]). The advantages here are the permanent availability of LLMs, their scalability and their ability to generate natural language texts, which are often indistinguishable from human-like texts [8]. Nevertheless, these huge language models were trained on a very large amount of data (often several trillion tokens). These data contain the social bias that we humans have, which is adopted by the LLM. In addition, these language models tend to hallucinate [9]. In the context of feedback generation this raises critical ethical considerations and demands rigorous assessment of effectiveness. One way to reduce the risk of AI based feedback is the integration of multiple feedback options as can be seen in figure 1.

In our Intelligent Tutoring System (ITS) the novice counselor writes with an AI based virtual client that has a psychosocial problem and the counselor gets feedback afterwards as can be seen in figure 1. An ITS is a platform that combines AI strategies with educational methodologies to create adaptive learning environments [10]. A detailed description of the ITS architecture and the user flow can be found in [11].

In this paper we investigate the nuanced roles that LLMs can

play in digital online counseling feedback systems, examine the capacity of LLMs to provide empathetic, objective, and appropriately lengthy feedback, compare various feedback methods, and evaluate the emotional and objective impacts of different LLMs. Additionally, we explore how AI feedback compares with peer feedback and discuss the potential risks associated with AI in this context as well as mitigation strategies. Overall, we address the following research questions:

- How can LLMs provide suitable feedback on an emotional and objective level in text-based counseling sessions?
- What is the optimal length for LLM-generated feedback in online counseling?
- How can the performance of LLMs (in terms of empathy, objectivity, and feedback length) be enhanced through fine-tuning?
- How is AI-based feedback accepted in society and especially among prospective online counselors?
- What are the potential risks associated with AI-based feedback in online counseling, and how can these be minimized?
- How does AI-based feedback compare to peer feedback in training scenarios for online counselors?

This structured approach to AI feedback within the domain of online counseling aims to contribute insights into the enhancement of digital counseling practices and the ethical deployment of AI technologies.

II. RELATED WORK

The history of ITS dates back to 1970 when Jamie Carbonell designed a program called SCHOLAR [12], which is often referred to as the first ITS [13], [14]. SCHOLAR utilized natural language to respond to a learner's question, pose a question, and provide feedback on the accuracy of the learner's answer. Since then, many other ITSs have been developed, leveraging various AI techniques such as Bayesian methods, NLP-based machine learning classifiers, and fuzzy-based techniques [10].

AI-based feedback systems have demonstrated potential across various domains, including programming [15], [16], driving behavior [17], mathematics [4], [18], [19], electronics [20], health science [21] and machine learning/data science [3], [22]. For instance, McDonald et al. [21] show that an AI tutor based on a dialogue manager (simple finite-state architecture) can already provide individualized feedback to students in large undergraduate classes in health science.

Due to the improvement by LLMs in the last two years, some research has focused on the evaluation of GPT3.5 or GPT4 models as tutors. Cao shows that combining GPT-3 and story-based gamification can support the learning of programming languages and increase the sense of belonging of Chinese students in introductory courses [23]. Dai et al. conducted a case study comparing feedback from ChatGPT with feedback from the instructor of a postgraduate course teaching introductory data science [3]. Similarly, Wang et al. explored the potential of large language models (LLMs) to

bridge the novice-expert knowledge gap in addressing math mistakes [4]. Chiu et al. [24] developed a computational framework for assessing LLM therapists' feedback by scoring generated responses and identifying areas for improvement.

Chaszczewicz et al. proposed a multi-level feedback generation approach using LLMs to provide feedback to novice counselors [7]. They developed a framework that checks the counselor's response to ensure it matches the counseling objective, automatically identifies areas for improvement (e.g., reflection or empathy), and suggests alternative goal-oriented responses.

While [7] focuses on providing real-time feedback to counselors during sessions, our approach emphasizes generating feedback after the counseling session. Additionally, our method integrates feedback more broadly into an ITS, expanding beyond the AI aspect alone to enhance the overall feedback mechanism.

III. ETHICAL CONSIDERATIONS

As Hatti and Timperley have stated, feedback is "one of the most powerful influences on learning and achievement, but this impact can be either positive or negative." [1] This quote underscores that feedback can not only enhance learning but also potentially decrease motivation if not handled correctly.

Glickman and Sharot [25] demonstrate that bias in AI can significantly influence human perceptions, emotions, and social judgments through feedback loops in human-AI interactions. Their research reveals that AI systems, by amplifying existing biases present in training data, can induce greater biases in humans who interact with them.

Concerns about deploying automated feedback can be divided into the five dimensions of trustworthy AI, named acceptance, explainability, accountability, fairness and privacy [26], [27]. While this paper primarily focuses on the acceptance of automated feedback generation in ITS, it also addresses ways to mitigate risks associated with the use of LLMs and improve understanding.

A. Acceptance

User acceptance of an AI system refers to their willingness to use it during interactions with a service [28]. To explore this, we conducted two surveys on feedback acceptance. The first survey addressed automated feedback generation in general, while the second concentrated specifically on its application in online counselor training with the presented ITS approach.

a) General Survey: The general survey, encompassing 71 participants from diverse backgrounds, including varying ages, professions, and other demographics, aimed to explore various aspects of feedback acceptance, including attitudes towards AI-generated feedback. Participants were asked about their familiarity with feedback mechanisms, the importance attributed to feedback, and their willingness to accept feedback from AI systems.

The results of the survey are illustrated in Figure 2. Chart a) represents the entire participant group, chart b) depicts the




Fig. 2. Willingness to accept feedback from AI Model

willingness to accept AI-generated feedback among participants with prior experience using LLMs, and chart c) displays this willingness among participants without such experience.

Although a significant correlation between experience and willingness to accept AI-generated feedback was not observed, there is a positive correlation ($r = 0.19$) present. This indicates a potential trend where increased experience may be associated with a greater willingness to accept AI-generated feedback, suggesting a possible influence of familiarity and exposure to AI technologies on individuals' attitudes towards feedback acceptance.

Acknowledging the limitations inherent in the participant sample, including its heterogeneous composition, the survey provides valuable insights into the complex interplay between experience and attitudes towards AI-driven feedback acceptance, with implications for future research and application in this domain.

b) Targetgroup Survey: Since the ITS is intended for use in the education of novice counselors, we focused on students of social work as our target group. We conducted another quantitative survey incorporating parts of the technology acceptance model (TAM) by Venkathesh and Davis [29], which helps to understand why people use or do not use a technology. This survey is conducted using an unmoderated remote test in conjunction with the ITS. Participants followed a schedule divided into three thematic blocks, allowing them to test the prototype and feedback options directly. GPT-4-1106-preview was used for feedback generation. GPT-4 was prompted with the prompt described in section IV, whereas no specific feedback method was applied here.

After interacting with the chatbot and receiving AI-generated feedback, participants provided their opinions on the feedback's characteristics. The survey included social work students, with a net response rate averaging over 90% and a total of 41 participants. Most students (89%) agreed or strongly agreed that the AI-generated feedback was justified, and they rated it as highly beneficial. Additionally, 96% received specific tips on how to improve their interactions, and they found the feedback length appropriate. These initial results indicate a positive assessment of AI feedback by the students.

B. Other dimensions of trustworthy AI

a) Explainability: In the context of automatic feedback generation, explainability means that users can understand how the AI generates feedback. Since auto-regressive LLMs learn from pre-training with large amounts of data (e.g., 2 trillion tokens with Llama2 [30] which is the base model of the Vicuna Model) they often function as "black-boxes" with unclear internal mechanisms. However, there are several approaches to make the LLM more applicable. Approaches to Explainable AI (XAI) include Local Analysis, which explains specific predictions, and Global Analysis, which aims to understand the model's overall knowledge and patterns [31]. Enhancing transparency in ITS development, such as showing system prompts and explaining the context, can also improve explainability. Additionally, it is important to note that the applicability of XAI algorithms is related to the openness of the model. In the case of highly intransparent models such as GPT3.5 or GPT4, not all XAI algorithms are applicable.

b) Accountability: Accountability refers to the responsibility of AI systems for their decisions. Given that state-of-the-art AI algorithms like LLMs cannot always generate accurate feedback, incorporating mechanisms for AI accountability is crucial to maintain stakeholders' trust [26]. For ITS development, it is recommended integrating a feedback mechanism that allows users to report issues to the development team.

c) Fairness: Fairness means that the users of the ITS should be treated equally and that the AI should not discriminate specific groups or individuals based on their counseling [32]. To detect a social bias or hurtful behavior of the LLM, our ITS is designed to display the AI-based feedback to the trainer as well. The trainer also has the opportunity to provide feedback to the development team at any time. Another option is to have the trainer confirm the feedback generated by the LLM and indicate this to the novice counselor with a badge. Effective prompt design is crucial in this process as it helps guide the AI in generating appropriate responses. Well-crafted feedback prompts might minimize the risk of unintended bias by ensuring that the AI's outputs are aligned with the rules of effective feedback which are described in section IV-A.

d) Privacy: Privacy ensures that sensitive information shared by individuals or gathered by AI systems is shielded from unauthorized or unlawful collection and usage [27]. We ensure privacy through a safe login mechanism, the mandatory entry of pseudonyms when registering for a course on our

ITS, so no peers and trainers can see the real names with the exceptions of examination results, where trainers can view the real name for grading purposes. Another important option to ensure privacy is to use and host an open-source LLM on premise instead of relying on AI providers like OpenAI or Anthropic when using an ITS for automatic feedback generation in production.

IV. FEEDBACK METHODS AND PROMPTING

In learning contexts, the way in which feedback is given plays a crucial role in enhancing relationships and skills, addressing mistakes, and improving behavior [33]. The following sections will discuss various established methods for giving and receiving effective feedback.

A. Principles for Effective Feedback

Effective feedback includes using descriptive language, being specific and behavior-related, and approaching feedback clearly and supportively [34, p. 24]. Key principles include:

- Focusing on specific observations.
- Avoiding vague statements.
- Using positive language.
- Concentrating on behaviors, not personal traits.
- Making feedback voluntary.
- Ensuring timely feedback.
- Being clear and concise.

De Villiers' [35] seven principles for effective feedback support and complement Fenger's foundations. These principles emphasize the need for situational, specific, meaningful, timely, relevant, and reliable feedback. Combining principles from various sources provides a comprehensive foundation for effective feedback design.

In addition to these principles, the usage of "I-messages" is crucial in feedback contexts, as it enables the feedback giver to express personal observations and intentions clearly [36, p. 6]. Utilizing I-messages helps maintain the subjectivity of the feedback, promoting openness, particularly in scenarios requiring critical feedback. In contrast, generalized "you-messages" can be perceived negatively, providing limited contextual insight and reducing the clarity of the intended feedback. Specific and clear I-messages enhance understanding and contribute to constructive feedback communication.

Many individuals hesitate to give "negative" feedback as criticism can cause discomfort and lead to avoidance [37]. However, in a learning platform aimed at skill development, addressing mistakes and improvement opportunities is crucial. It is important to differentiate between personal and professional critiques. Personal criticism often triggers emotional responses and is harder to accept, whereas professional feedback should be seen as valuable guidance for improvement. Thus, negative feedback should be actively sought and viewed as a path to personal growth. Applying specific feedback methods can be beneficial. The following section introduces three effective and constructive feedback methods and demonstrates how these methods were used in prompts: the Sandwich Method, the WWW Method, and the STATE Method.

B. Prompting

1) *Basic Prompt Structure:* To develop the evaluation application for AI-generated feedback, prompt engineering was conducted to instruct large language models (LLMs) to provide effective feedback on counseling texts. The main challenge was delivering sufficient context without overloading the models. The final prompt structure evolved through various iterations, ensuring clarity and conciseness. The best results were achieved by providing instructions in short, clear sentences. Originally, these prompts were formulated in German, but they have been translated into English for this paper. Below is the fundamental framework of the final prompt for feedback from a mentor's perspective:

Prompt template of the AI mentor

Act as a mentor providing constructive and learning-oriented feedback. You receive a counseling chat between the client {client} and their social counselor (user). To improve the counselor's skills, it is important that they receive quality feedback. Your feedback should be descriptive rather than evaluative, specific to behaviors, and inviting. Frame your feedback using I-messages. Please keep your feedback brief and to the point. Adhere to the following feedback method: {method}. Description of the method: {method description}. Do not use greetings or farewells. Provide specific improvement suggestions if necessary. Chat transcript: {Chathistory}. Now give the counselor feedback from the mentor's perspective.

- **Role Introduction:** "Act as a mentor providing constructive and learning-oriented feedback. You receive a counseling chat between a client {client} and their counselor (user). To improve the counselor's skills, it's essential they receive quality feedback."
- **Feedback Nature:** "Your feedback should be descriptive rather than evaluative, specific to behaviors, and inviting."
- **Use of I-Messages:** "Frame your feedback using I-messages to maintain a personal tone."
- **Conciseness:** "Keep your feedback brief and to the point."
- **Method Specification:** "Adhere to the following feedback method: {method}. Description: {method description}."
- **Avoid Formalities:** "Do not use greetings or farewells."
- **Concrete Suggestions:** "Provide specific improvement suggestions if necessary."
- **Role Reminder:** "Provide specific improvement suggestions if necessary."

2) *Feedback Methods Prompt Structure:* The following sections describe these feedback methods in detail and provide their corresponding prompt templates. These feedback method prompt descriptions are inserted into the {method description} placeholder in the main prompt above.

Sandwich Method: The Sandwich Method involves placing a criticism (“filling”) between two positive comments (“bread”) [38, p. 68f]. The first positive comment highlights specific positive observations, while the last emphasizes a general positive aspect. This method is considered quick and efficient [39, p. 43ff]. However, it is crucial that the positive comments are genuine and not just a means to deliver criticism. For example, statements like “I generally liked it, but...” are not helpful. Instead, positive comments should reinforce desired behavior. Studies have shown that the order of Positive-Criticism-Positive is not always the most effective. An alternative order, Criticism-Positive-Positive, has been found to be more impactful as it reduces the amount of context before corrective feedback [40].

Sandwich Method Prompt Description

Criticism is sandwiched between 2 positive aspects. The first positive point is a specific positive observation. Then comes the criticism. If there isn’t any, say so. The third point is another general positive aspect. Keep it brief and focus on the essentials.

WWW Method: The WWW feedback method provides a structured approach with three components: Perception (German *Wahrnehmung*), Impact (German *Wirkung*), and Wish (German *Wunsch*) [41, p. 60]. This structure facilitates the identification of strengths and weaknesses by detailing how behaviors are perceived and their effects. Initially, an observation is outlined (“I noticed that...”), followed by its impact (“This makes me feel...”), and concluding with a wish or suggestion (“In the future, I would like...”). A common difficulty in employing this method is separating perception and impact accurately. For instance, instead of stating “I noticed you were not empathetic during the session,” a clearer expression would be “I noticed that your responses did not sufficiently address the emotional needs of the person.”

WWW Method Prompt Description

Give your feedback in these categories: Perception, Impact, and Wish. Perception = Concrete description of observations. Impact = Description of how the situation or behavior affected you. This works for both positive and negative feedback. Wish = What could be done better in the future. Keep it short and concise. Provide a very brief summary at the end. Structure your feedback in 3 steps: I noticed that you..., My impression was that you..., I would appreciate it if you..

STATE Method: The STATE Method expands on the WWW Method by adding additional aspects [42]. STATE stands for:

- Share the facts: Express your perception.
- Tell your story: Describe the impact.

- Ask for others’ paths: Invite the other person to share their perspective.
- Talk tentatively: Phrase your feedback tentatively.
- Encourage testing: Encourage the other person to express opposing views.

The first two aspects mirror the WWW method, but the STATE method initiates a dialogue by asking the other person for their perspective (“Can you explain why you responded that way?”). The last two points focus on how to deliver this feedback, encouraging open and tentative communication. This method goes deeper by actively involving the other person in the feedback process, which may not be feasible in all contexts.

STATE Method Prompt Description

The term STATE consists of the English terms of the method: Share the facts: Express perception. Tell your story: Communicate the impact (of perception). Ask for others’ paths: Ask the other person about their perspective. Talk tentatively: Formulate tentatively. Encourage testing: Encourage the other person to express opposing views. Structure your feedback in these steps: I noticed that you..., I wondered if you did that because..., Can you tell me why you did ...?, Can you tell me what I might have overlooked?

These descriptions of the feedback methodology also underwent several iterations. It was crucial to describe the methods as briefly as possible while still being clear and explicit. For the WWW and STATE methods, it was necessary to provide a basic framework for formulation, as the responses were previously varied in structure.

Additionally, the models were prompted to generate a fourth feedback, this time without a predefined methodology. This approach aimed to ascertain whether the models are capable of producing well-structured feedback even in the absence of clear guidelines.

No Method Prompt Description

No specific method prescribed. Provide feedback freely. Keep it brief and focus on the essentials.

By incorporating these feedback methods into the prompting structure, the AI models were guided to provide feedback that aligns with established principles and methodologies, ensuring consistency and effectiveness in the feedback process.

In addition to feedback from a mentor’s perspective with these methods, feedback was also generated from the viewpoint of potential clients. This was achieved using a standardized questionnaire format specifically designed to capture the satisfaction and perceived quality of the counseling from their perspective. This questionnaire is part of the Integrative Quality Assurance Model (IQSM) developed by Eidenbenz and Lang, as detailed in [43, p. 220ff]. This integration ensures that not only the perspective of the counselor is considered,

but also direct feedback from potential clients is incorporated to comprehensively evaluate the quality of the counseling.

Simulated Client Prompt Description

You receive a chat transcript between the client {client} and their social counselor. To help the counselor improve their skills, feedback should be generated from the client's perspective. This will be done in the form of a questionnaire: Rating 1 for very, 2 for mostly, 3 for only partially, 4 for little: I am generally satisfied with the counseling. The counselor understood my concern. The counselor took my question seriously. The counseling was helpful in clarifying my concern. The answer provided me with a new perspective. I could implement the insights into my practice. The counselor chose the right words and the right tone. I would contact this counselor again. Free text feedback: Additional comments. You should respond in the following format and insert your ratings: { "satisfied": , [...] "contact_again": , "free_text_feedback": , } Please also provide specific feedback in "free_text_feedback" about what you were particularly satisfied with or not satisfied with. The counselor should understand what could be improved. [...] Chat transcript: Chathistory. Now provide the counselor feedback from the client's perspective.

The following is an example of free text feedback, generated from the GPT-4 model from a client perspective, demonstrating the detailed and personal nature of the responses generated:

Sample Feedback from Client Perspective

I am very satisfied with the counseling. The counselor Maria understood me and my concerns very well and took them seriously. I felt supported and not left alone. Her questions were empathetic and helped me better understand my situation. The tips, such as leaving the room, are practical and implementable, but I might have wished for even more concrete assistance on how to strengthen my ability to say no. I greatly appreciate that a follow-up appointment was arranged, and Maria struck the right tone to create a trusting atmosphere. I will definitely reach out to her again.

V. METHODS

A. Evaluation of AI based Feedback

To systematically verify the accuracy and quality of AI-generated feedback an evaluation tool has been developed. The feedback was generated for chat counseling sessions which were created in student's role plays in a course for Online Counseling. While one student received a description for a client with a psycho-social problem in the context of addiction or family counseling, another student was given the task of conducting an initial counseling session with this client. To get

a variety of different counseling sessions seven different client descriptions were created. All in all the dataset contains 64 German counseling dialogues with an average of 38 messages per session. The dataset was supplemented with feedback generated by four models:

- GPT-3.5-turbo-1106 by OpenAI
- GPT-4-1106-preview by OpenAI
- Vicuna-13b-v1.5-16k by LMSYS Org
- Mixtral-8x7b1 by MistralAI

The models were instructed to generate feedback from a mentor's perspective using the Sandwich, WWW, and STATE methods, as well as an additional feedback without a pre-defined method. Additionally, feedback from the perspective of the client was generated using a questionnaire format. Thus, each conversation received 20 AI-generated feedbacks for evaluation. At the end of the evaluation process, a total of 1280 feedback instances were available for analysis (4 models × 5 feedback methods × 64 conversation transcripts). This comprehensive dataset provided a robust basis for assessing the effectiveness and adherence to feedback principles across different AI models and methods.

In the evaluation application, three levels were used to assess satisfaction with AI-generated feedback. Participants rated feedback on:

- Emotional Level: How would you feel receiving this feedback? (Scale from 0 - Very Poor to 4 - Excellent)
- Objective Level: How suitable do you find the feedback content according to the described method? (Scale from 0 - Very Inappropriate to 4 - Very Appropriate)
- Feedback Length: What is your opinion on the length of the feedback text? (Options: Too Short, Just Right, Too Long)

Which means 4.0 represents the highest achievable rating, while 0 denotes the lowest. Three raters conducted the evaluation. Those three raters are student assistants that have experience in online counseling methods and have received an expert briefing on the topic of giving feedback. The evaluation with the raters was carried out twice, whereby in the second evaluation none of the raters rated the same feedback again. Afterwards the average values of both runs were taken for evaluation. The evaluation was also accompanied by experts who randomly checked the feedback assessment. Approximately 70% (840) of the 1280 feedback instances were assessed in two iterations. To publish LLMs capable of providing high-quality feedback, we propose to fine-tune the smallest tested LLM (Vicuna-13B-16K) using the highest-rated feedback from the evaluation. Using the smallest open-source LLM (Vicuna-13B-16K) for fine-tuning is advantageous due to its resource efficiency, requiring less computational power and memory, which makes the process cost-effective and faster. Additionally, open-source models offer accessibility, allowing broader experimentation and adaptation without restrictive licensing. Furthermore, smaller models are easier to deploy across various platforms, ensuring wide usability. By fine-tuning with the highest-rated feedback, we try to ensure high-

quality output while maintaining manageability and scalability, ultimately leading to a practical and effective implementation of high-quality feedback systems. The remaining 30% of the feedback instances were subsequently utilized for human evaluation of the fine-tuned model, which was refined with the best-rated feedback during this evaluation. This process is detailed in the next section.

B. Fine-Tuning

As shown in Section VI-A, the Vicuna-13b-v1.5-16k model performs poorly compared to larger models. Consequently, we used the highest-rated feedbacks of the described feedback evaluation to create a dataset for LoRa-fine-tuning the Vicuna-13b-v1.5-16k model. To be included in the fine-tuning dataset, feedback must score at least six points in total when summing the emotional and content level evaluations. The training parameters are detailed in the appendix, Table II. All in all 505 feedbacks were used for fine-tuning and 170 (19 conversation transcripts x 5 feedback methods x 2 methods)

C. Comparison of Peer-Feedback and AI based Feedback

As described in our study on functionality and acceptance (paragraph III-A0b), we also wanted to compare the AI-generated feedback with feedback from human peers. Students were tasked with providing self-written feedback to a fellow student based on the recipient's conversation history. Additionally, participants received content-related guidelines for feedback, as detailed in section IV-A. The students were asked to use between 120 and 180 words to express their feedback. Samples of the human feedback were then analyzed by an expert in the field of online counseling. The study investigates the quality and effectiveness of feedback provided in response to counseling conversations, systematically comparing 13 instances of AI-generated feedback to peer-generated feedback.

VI. RESULTS

A. Results of AI based Feedback Evaluation

In the overall assessment, it is notable that the models generally achieved higher scores in objective evaluation (average total score of 2.92) compared to emotional evaluation (average total score of 2.56), as detailed in Figure 3 on the left side. This trend suggests that the models excel in meeting objective criteria, while emotional aspects may not be addressed to the same extent. The GPT-4 model stands out with an average objective rating of 3.54 and an emotional rating of 2.87. In contrast, the Vicuna model performed poorest overall, with an average objective rating of 2.02 and an emotional rating of 2.10, making it the only model where emotional evaluation slightly surpassed the objective score. However, this slight difference does not significantly alter the overall trend, as the Vicuna model still achieved the lowest performance in both categories.

To provide concrete examples of how these variations are apparent in practical scenarios, we present excerpts of feedback generated by GPT-4 and Vicuna for the same chat history with the "No Method" approach:

Sample Feedback from Vicuna

1. You are very open and inviting in your communication, which sends a positive signal to the client.
2. You are very concrete and vivid in describing situations and emotions, which helps the client better understand what they are going through.
3. You ask open-ended questions to encourage the client to reflect and to give them the opportunity to organize their thoughts and feelings.
4. You offer concrete solutions and tips to help the client cope with their problems.
5. You provide the client with space to talk about their feelings and thoughts by allowing them to freely share their opinions and thoughts.
6. You are very attentive to the client's needs and work closely with them to find solutions that are best suited for them.

Sample Feedback from GPT-4

I have carefully read the chat history between you and Jessica and would like to give you my feedback. First of all, I would like to highlight positively that you chose a friendly and open approach to start the conversation. [...] Your questions are open-ended, which encourages Jessica to talk about her situation. You have also acknowledged her feelings and concerns, creating a trusting atmosphere. [...] One area where you could possibly improve is to show empathetic concern earlier in the conversation when it comes to the topic of drinking behavior and alcohol consumption, and gently point out potential risks to Jessica. This could raise her awareness of the problem without confronting her too much. Furthermore, when it comes to creating a study plan or finding solutions for her stress situation, you could work with Jessica to develop concrete steps together, instead of giving suggestions. This would strengthen her autonomy and sense of responsibility. In conclusion, I find it very good that you offer a follow-up appointment, thus ensuring continuity in support. This gives Jessica the feeling of not being alone and being supported. Overall, you have shown a supportive and encouraging counseling attitude that is certainly helpful for Jessica. Keep it up!

Examining the differences in model performance with specific feedback methods provides insights into how structured approaches can influence feedback quality across both emotional and objective evaluations. This analysis is crucial for understanding which methods enhance or hinder the feedback generated by different models. These findings are visually represented in the heatmaps in Figure 3, with the middle heatmap showing emotional evaluation differences and the right heatmap showing objective evaluation differences.

The GPT-4 model is consistent across various methods but underperforms with the STATE method in both emotional and objective evaluations (0.5 below GPT-4 emotional rating av-




Fig. 3. Average ratings per model (left) and model performance per method (middle and right). Each column in the heatmaps contains the deviations per method to the average rating shown left for a specific model. Positive values indicate that a method works relatively well for the model.

erage of 2.87 and 0.3 below objective rating average of 3.54). In both emotional and objective evaluations, GPT-3.5 and Mistral-small perform best without a specific method (between 0.28 and 0.46 above respective model’s average), indicating their strength in generating feedback without constraints. Vicuna-13b consistently performs poorly, with the client and STATE method causing the largest negative deviation from its average score.

Specifically, the GPT-3.5 model significantly underperforms with the Sandwich method in the objective evaluation, showing a decrease of -0.53. The Mistral-small model experiences the highest overall deterioration with the client method, showing a decline of -1.1 in the emotional evaluation and -0.49 in the objective evaluation. Structured methods like Sandwich and WWW generally enhance performance across most models, reflecting their utility in promoting clear and constructive feedback. However, the STATE method often results in lower scores, suggesting it may be less effective.

Overall, while certain feedback methods can improve feedback quality, as shown in the heatmaps in Figure 3, the effectiveness of each method varies significantly across models. Understanding these nuances is essential for optimizing feedback methodologies to suit each model’s strengths.

B. Evaluation of Feedback Length

In Figure 4, the distribution of rated feedback lengths is illustrated. For the “Too long” category, the distribution is more spread out, with a higher concentration around mid-lengths (approximately 200-400 words). The “slightly too long” category, where the first and second epochs show disagreements between “Too long” and “Perfect Length” displays a higher density around 100-200 words but is more dispersed compared to other categories.

The “perfect length” length category has a more pronounced peak around 100-190 words, indicating a higher concentration of feedback perceived as the perfect length within this range. The “slightly too short” category exhibits a smaller range, with most feedback lengths clustering around 100 words or less. As expected, there is some overlap between these classes, reflecting the variability in feedback length depending on various factors. These factors include the feedback giver and




Fig. 4. Distribution of feedback text length ratings

receiver, the complexity and length of the counseling, as well as the quality of the counseling (e.g., criticism).

The distribution of length ratings, as depicted in Table I, offers insights into the perceived adequacy of generated content across different models. Notably, Vicuna-13b garnered the highest count of “too short” and “slightly too short” ratings, suggesting a tendency towards generating shorter responses compared to other models. Conversely, gpt-4 received the most “too long” and “slightly too long” ratings, indicating a propensity for verbosity in its outputs. Interestingly, despite variations in performance across other metrics, both gpt-3.5 and mistral-small received only one “too long” rating each, hinting at potential consistency in their generation of content length. The predominance of “perfect length” ratings across all models suggests a general satisfaction with the length of generated feedbacks.

TABLE I
FREQUENCY OF LENGTH RATINGS FOR DIFFERENT MODELS

length rating	GPT-3.5	GPT-4	Mistral	Vicuna	total
too short	6	1	18	39	64
slightly too short	51	1	31	45	128
exactly right	150	144	150	103	547
slightly too long	2	41	10	11	64
too long	1	23	1	12	37

When examining the correlation between feedback length rating and emotional and objective ratings, several patterns emerge, as shown in figure 5. Feedback perceived as “too short” consistently receives lower emotional (1.58) and objective (1.34) ratings. In contrast, feedback rated as “perfect length” scores the highest for both emotional (2.71) and objective (3.16) aspects. This suggests that there is an optimal range for feedback length, where it is considered thorough yet concise enough to be effective.

Interestingly, while “too long” feedback tends to score lower than “perfect length” feedback, it still fares better than “too short” feedback. This indicates that verbosity is less detrimental to perceived feedback quality than brevity. Thus, while concise feedback is ideal, slightly longer feedback is preferable to overly brief comments. Additionally, the higher ratings for longer feedback can be attributed to GPT-4, which generally provides longer feedback but consistently receives good ratings, as previously established.




Fig. 5. Correlation between Length Evaluation and Average Rating

C. Fine-Tuning

Figure 7 in the appendix illustrates the evaluation and training loss throughout the fine-tuning process of the Vicuna-13b-16k model. The model demonstrates improvement on the evaluation dataset up to step 88, after which the evaluation loss begins to rise. This increase may suggest overfitting to the training data, as the training loss continues to decrease until the end. Consequently, we used a checkpoint at step 122, since no checkpoint was made at step 88.

The bar chart in figure 6 compares the emotional and objective scores of the Vicuna model and its fine-tuned version across different categories: ‘No Method’, ‘Client’, ‘Sandwich’, ‘STATE’, and ‘WWW’. It also includes an ‘Average’ category to summarize overall performance. For each category pair, the emotional scores are represented by the bars on the left, while the objective scores are represented by the bars on the right. The changes due to fine-tuning are overlaid with either green (indicating improvement) or red (indicating deterioration) bars.

The Vicuna model received an average emotional rating of 2.37 and an objective rating of 2.6. In contrast, the fine-tuned model (vicuna-ft) achieved higher ratings in both dimensions, with an emotional rating of 2.6 and an objective rating of around 3.29. This indicates that the fine-tuned model outperforms the original Vicuna model in both emotionality




Fig. 6. Comparison of Emotional and Objective ratings between Vicuna and Fine-Tuned Vicuna across different feedback methods

(improvement of 0.23) and objectivity (improvement of 0.69). The gap between the improvements of emotionality and objectivity could be through the training data, since the training data had also higher ratings for objectivity. It shows that the fine-tuned version consistently outperforms the Vicuna model in both emotional and objective ratings, with the exception of emotionality for the category “Client”.

According to sample checks, the feedback from the fine-tuned Vicuna variant seems to contain more continuous text and focuses more often on the essentials of the counseling. Nevertheless, mistakes are often made and feedback from GPT-4, for example, still seems to provide feedback that is much more conducive to learning. An example of feedback from the fine-tuned Vicuna model is shown below:

Sample Feedback from fine-tuned Vicuna

I noticed that you created an open and warm atmosphere at the beginning of the conversation by giving Luisa the opportunity to speak to you in the way she wanted. You also quickly made a connection with Luisa by asking about her experience and making her feel that you understood her. My impression was that you listened well and encouraged Luisa to share her feelings. You also tried to encourage Luisa to have the conversation with her parents and helped her to develop a strategy to do this. I would appreciate it if in the future you could be even more responsive to Luisa’s needs and feelings and help her to develop concrete steps to improve her situation. It would also be helpful if you could encourage Luisa to recognize and use her own strengths and resources to cope with her situation. Overall, you did a good job of supporting Luisa and making her feel understood and taken seriously.

D. Comparison of Peer-Feedback and AI based Feedback

The analysis of the sample feedbacks shows that all feedback used positive and appreciative language. A noticeable difference is the length of the feedback: peer feedback was significantly shorter and often more general and vague. For instance, general statements like “...I found your chat very

helpful... " were common, with little specific observation. This could be due to the short processing time of the feedback providers and the nature of the test task.

In some cases, peer feedback included specific comments based on the conversation context, but these were often only related to the initial phase of the chat counseling. These comments were positive but not very detailed: "... *your counseling was well done. I especially liked your introduction with the framework conditions.*"

In contrast, the AI-generated feedback was more detailed and specific. It highlighted the importance of the framework conditions and provided constructive criticism and suggestions for improvement: "*Additionally, you asked Lars what state he wants to achieve in 3-4 weeks, which shows forward-thinking. It would have been helpful if you had also helped him set smaller, achievable goals to support him on his journey and give him a sense of progress.*"

The AI feedback also included concrete examples that could be applied in future exercises: "*For example, you could say: 'It sounds like you're going through a really challenging time. Would you like to tell me more about how you feel?'*"

The AI feedback was structured and presented clearly according to the given prompts, which can be a double-edged sword. Highly motivated students might get demotivated by constantly receiving new improvement suggestions. Therefore, feedback should come from various sources, including human feedback from course instructors and other students. This provides a mix of professional feedback and valuable perspectives from peers.

The quality of peer feedback could be improved through targeted exercises for the students. Additionally, students should be able to critically reflect on the AI feedback to avoid biases in their perceptions. A significant advantage of AI-generated feedback is that it can be requested promptly after an exercise, providing immediate feedback and eliminating the waiting time.

VII. LIMITATIONS AND FUTURE WORK

While this study provides valuable insights into the implementation and evaluation of AI-generated feedback in counselor education, several limitations should be acknowledged. These limitations also suggest directions for future research and practice.

It is important to acknowledge that GPT-3.5 and GPT-4 are closed-source language models, raising ethical concerns about their use in automatic feedback generation. Furthermore, XAI approaches cannot be applied to closed-source LLMs, and there is no guarantee regarding the fate of the data involved.

The study's findings are based on a sample of 64 conversations in the context of addiction and family counseling, which, although valuable, may not fully represent the wide variety of scenarios encountered in counselor education. Future research with larger and more diverse datasets could help validate these findings and enhance their applicability across different educational contexts and counseling situations. The analysis comparing AI-generated feedback to peer feedback involved

only 13 instances, which may not capture the full spectrum of feedback quality and styles. Future research with larger and more diverse datasets could help validate these findings and enhance their applicability across different educational contexts and counseling situations.

Our evaluation focused primarily on emotional and objective criteria, as well as feedback length. While these metrics offer important insights on feedback quality, they do not capture all dimensions of effective feedback, such as specificity of improvement suggestions. Expanding the range of evaluation metrics in future studies could provide a more comprehensive understanding of AI-generated feedback's impact.

The immediate reception and perceived quality of AI-generated feedback were the primary focus of this study. Future research should explore the long-term effects of such feedback on learners' skill development and overall learning outcomes. Longitudinal studies tracking progress over time would offer valuable insights into the sustained benefits and potential limitations of integrating AI feedback into counselor education.

This study utilized specific AI models, including GPT-3.5, GPT-4, Vicuna-13b, and Mixtral-8x7b1. While these models represent a range of capabilities, the findings may not be directly applicable to other AI models or future iterations. Broadening the scope to include a wider array of models in future research will help ensure the findings' relevance as AI technology evolves.

Future work should also delve deeper into the development of feedback-mechanisms for specific counseling methods and feedback at the utterance level. This could enhance the granularity and relevance of AI-generated feedback. Additionally, research should focus on longitudinal studies to assess the long-term impact of AI feedback on counselor education. Exploring ways to integrate AI feedback seamlessly with human feedback will also be crucial to develop a balanced and effective training ecosystem. Furthermore, addressing ethical considerations in greater depth, will be essential as AI becomes more integrated into educational settings.

By acknowledging these limitations and outlining directions for future research, we aim to provide a pathway for continued improvement in the effectiveness and ethical implementation of AI-generated feedback in counselor education.

VIII. CONCLUSION

This study explored the implementation and evaluation of AI-generated feedback in an ITS designed for counselor education. The findings demonstrate that AI models can generate detailed and specific feedback that meets objective criteria effectively. However, challenges remain in ensuring that the feedback resonates emotionally with learners, especially when using smaller open source models. The research questions described in section I are answered concisely below.

AI models can indeed provide suitable feedback on both an emotional and objective level in text-based counseling sessions, though emotional resonance remains an area for

improvement. The optimal length for LLM-generated feedback tends to be between 100-190 words. Feedback that is perceived as “too short” often receives lower emotional and objective ratings, whereas feedback considered “too long” is less detrimental but still less effective than feedback of optimal length.

Different feedback methods such as Sandwich, WWW, and STATE do not consistently improve the quality of AI-generated feedback. Some models perform better without the constraints of a specific feedback method, indicating that while structured approaches can offer clarity, they may not always enhance effectiveness. The study also found that AI-generated feedback tends to score higher in objective evaluation compared to emotional evaluation, suggesting a need for improvements in emotional resonance.

Fine-tuning can enhance the performance of LLMs, as evidenced by improvements in both emotional and objective ratings for the Vicuna model after fine-tuning. This indicates that targeted fine-tuning can effectively enhance the empathetic and objective aspects of feedback generation. AI-based feedback is generally accepted, with higher acceptance among individuals who have prior experience using AI systems. Surveys indicate that familiarity with AI technologies positively influences attitudes toward AI-generated feedback, suggesting that exposure and education can enhance acceptance.

AI-generated feedback was found to be more detailed and actionable compared to peer feedback, highlighting its potential to supplement traditional feedback mechanisms in counselor education. Despite the promising results, the study’s limitations include a relatively small and specific sample size, a narrow scope of evaluation metrics, a small number of survey participants, and the limited analysis of only 13 instances of peer feedback. These limitations suggest several directions for future research, including the need for larger and more diverse datasets, a broader range of evaluation metrics, deeper exploration of ethical issues, and longitudinal studies to assess the long-term impacts of AI feedback.

Potential risks associated with AI-based feedback include the propagation of social biases present in training data and the possibility of AI models “hallucinating” incorrect or misleading feedback. Mitigation strategies include integrating multiple feedback options, continuous evaluation and fine-tuning of models, and ensuring transparency and accountability in AI-based feedback systems. Additionally, the use of a mentored control instance, where feedback generated by AI is reviewed and moderated by a human expert, can help mitigate these risks by providing an additional layer of oversight and ensuring that the feedback remains accurate and contextually appropriate.

In conclusion, while AI-generated feedback offers a powerful tool for counselor education, its implementation must be approached with careful consideration of its limitations and ethical implications. By addressing these areas in future research, we can ensure that AI systems effectively support the development of counseling competencies and contribute positively to the educational experience.

TABLE II
PARAMETER SETTINGS FOR FINE-TUNING

Parameter	Value
Sequence Length	4096
Sample Packing	True
Pad to Sequence Length	True
Eval Sample Packing	False
Adapter Type	LoRA
LoRA Rank (r)	32
LoRA Alpha	16
LoRA Dropout	0.05
LoRA Target Linear	True
LoRA Fan In/Fan Out	Enabled
Gradient Accumulation Steps	4
Micro Batch Size	2
Number of Epochs	6
Optimizer	AdamW
Learning Rate Scheduler	Cosine
Learning Rate	0.0002
Flash Attention	True
Eval Max New Tokens	256

APPENDIX




Fig. 7. Training and Evaluation Loss during the fine-tuning process with 6 epochs

REFERENCES

- [1] J. Hattie and H. Timperley, “The Power of Feedback,” *Review of Educational Research*, vol. 77, no. 1, pp. 81–112, Mar. 2007. doi: 10.3102/003465430298487
- [2] S. Narciss, S. Sosnovsky, L. Schnaubert, E. Andrès, A. Eichelmann, G. Gogudze *et al.*, “Exploring feedback and student characteristics relevant for personalizing feedback strategies,” *Computers & Education*, vol. 71, pp. 56–76, Feb. 2014. doi: 10.1016/j.compedu.2013.09.011
- [3] W. Dai, J. Lin, H. Jin, T. Li, Y.-S. Tsai, D. Gašević *et al.*, “Can Large Language Models Provide Feedback to Students? A Case Study on ChatGPT,” in *2023 IEEE International Conference on Advanced Learning Technologies (ICALT)*, Jul. 2023. doi: 10.1109/ICALT58122.2023.00100 pp. 323–325, iSSN: 2161-377X.
- [4] R. E. Wang, Q. Zhang, C. Robinson, S. Loeb, and D. Demzky, “Bridging the Novice-Expert Gap via Models of Decision-Making: A Case Study on Remediating Math Mistakes,” Apr. 2024, arXiv:2310.10648 [cs].
- [5] E. M. Engelhardt, “Onlineberatung – Digitales Beratungsangebot für Alle?” in *Digital Diversity*, H. Angenent, B. Heidkamp, and D. Kergel, Eds. Wiesbaden: Springer Fachmedien Wiesbaden, 2019, pp. 161–173. ISBN 978-3-658-26752-0 978-3-658-26753-7
- [6] M. Stieler, S. Lipot, and R. Lehmann, “Zum Stand der Onlineberatung in Zeiten der Corona Krise. Entwicklungs- und Veränderungsprozesse der Onlineberatungslandschaft,” *e-beratungsjournal.net – Zeitschrift für Online-Beratung und computervermittelte Kommunikation*, vol. 18,

- no. 1, pp. 50–65, 2022. doi: 10.48341/262P-7T64 Publisher: Universität für Weiterbildung Krens & e-beratungsjournal.net.
- [7] A. Chaszczewicz, R. S. Shah, R. Louie, B. A. Arnou, R. Kraut, and D. Yang, “Multi-Level Feedback Generation with Large Language Models for Empowering Novice Peer Counselors,” Mar. 2024, arXiv:2403.15482 [cs].
- [8] S. Minaee, T. Mikolov, N. Nikzad, M. Chenaghlu, R. Socher, X. Amatriain *et al.*, “Large Language Models: A Survey,” Feb. 2024, arXiv:2402.06196 [cs].
- [9] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang *et al.*, “A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions,” Nov. 2023, arXiv:2311.05232 [cs].
- [10] E. Mousavinasab, N. Zarifsanaiy, S. R. Niakan Kalhori, M. Rakhshan, L. Keikha, and M. Ghazi Saeedi, “Intelligent tutoring systems: a systematic review of characteristics, applications, and evaluation methods,” *Interactive Learning Environments*, vol. 29, no. 1, pp. 142–163, Jan. 2021. doi: 10.1080/10494820.2018.1558257
- [11] E. Rudolph, N. Engert, and J. Albrecht, “An AI-Based Virtual Client for Educational Role-Playing in the Training of Online Counselors,” in *Proceedings of the 16th International Conference on Computer Supported Education - Volume 2: CSEDU*, vol. 2. SCITEPRESS, May 2024. doi: 10.5220/0012690700003693. ISBN 978-989-758-697-2 pp. 108–117.
- [12] J. Carbonell, “AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction,” *IEEE Transactions on Man-Machine Systems*, vol. 4, no. 11, pp. 190–202, 1970. doi: 10.1109/TMMS.1970.299942
- [13] A. T. Corbett, K. R. Koedinger, and J. R. Anderson, “Chapter 37 - Intelligent Tutoring Systems,” in *Handbook of Human-Computer Interaction (Second Edition)*, M. G. Helander, T. K. Landauer, and P. V. Prabhu, Eds. Amsterdam: North-Holland, Jan. 1997, pp. 849–874. ISBN 978-0-444-81862-1
- [14] W. Ma, O. O. Adesope, J. C. Nesbit, and Q. Liu, “Intelligent tutoring systems and learning outcomes: A meta-analysis,” *Journal of Educational Psychology*, vol. 106, no. 4, pp. 901–918, 2014. doi: 10.1037/a0037123 Place: US Publisher: American Psychological Association.
- [15] C. Cao, “Scaffolding CS1 Courses with a Large Language Model-Powered Intelligent Tutoring System,” in *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*, ser. IUI '23 Companion. New York, NY, USA: Association for Computing Machinery, 2023. doi: 10.1145/3581754.3584111. ISBN 9798400701078 pp. 229–232.
- [16] J.-Y. Kuo, H.-C. Lin, P.-F. Wang, and Z.-G. Nie, “A Feedback System Supporting Students Approaching a High-Level Programming Course,” *Applied Sciences*, vol. 12, no. 14, p. 7064, Jan. 2022. doi: 10.3390/app12147064 Number: 14 Publisher: Multidisciplinary Digital Publishing Institute.
- [17] Z. Marafie, K.-J. Lin, D. Wang, H. Lyu, Y. Liu, Y. Meng *et al.*, “AutoCoach: An Intelligent Driver Behavior Feedback Agent with Personality-Based Driver Models,” *Electronics*, vol. 10, no. 11, p. 1361, Jan. 2021. doi: 10.3390/electronics10111361 Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.
- [18] A. Botelho, S. Baral, J. A. Erickson, P. Benachamardi, and N. T. Hefernan, “Leveraging natural language processing to support automated assessment and feedback for student open responses in mathematics,” *Journal of Computer Assisted Learning*, vol. 39, no. 3, pp. 823–840, 2023. doi: 10.1111/jcal.12793
- [19] B. Grawemeyer, M. Mavrikis, W. Holmes, S. Gutierrez-Santos, M. Wiedmann, and N. Rummel, “Affecting off-task behaviour: how affect-aware feedback can improve student learning,” in *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, ser. LAK '16. New York, NY, USA: Association for Computing Machinery, Apr. 2016. doi: 10.1145/2883851.2883936. ISBN 978-1-4503-4190-5 pp. 104–113.
- [20] M. Dzikovska, N. Steinhilber, E. Farrow, J. Moore, and G. Campbell, “BEETLE II: Deep Natural Language Understanding and Automatic Feedback Generation for Intelligent Tutoring in Basic Electricity and Electronics,” *International Journal of Artificial Intelligence in Education*, vol. 24, no. 3, pp. 284–332, Sep. 2014. doi: 10.1007/s40593-014-0017-9
- [21] J. McDonald, A. Knott, S. Stein, and R. Zeng, “An empirically-based, tutorial dialogue system: design, implementation and evaluation in a first year health sciences course,” in *Proceedings of Electric Dreams. Proceedings ascilite 2013*. Australasian Society for Computers in Learning in Tertiary Education, 2013. ISBN 978-1-74138-403-1 pp. 562–572.
- [22] F. St-Hilaire, D. D. Vu, A. Frau, N. Burns, F. Faraji, J. Potochny *et al.*, “A New Era: Intelligent Tutoring Systems Will Transform Online Learning for Millions,” Mar. 2022.
- [23] C. Cao, “Leveraging Large Language Model and Story-Based Gamification in Intelligent Tutoring System to Scaffold Introductory Programming Courses: A Design-Based Research Study,” Feb. 2023, arXiv:2302.12834 [cs].
- [24] Y. Y. Chiu, A. Sharma, I. W. Lin, and T. Althoff, “A Computational Framework for Behavioral Assessment of LLM Therapists,” Jan. 2024, arXiv:2401.00820 [cs].
- [25] M. Glickman and T. Sharot, “How human-AI feedback loops alter human perceptual, emotional and social judgements,” Nov. 2022.
- [26] J. Lin, L. Sha, Y. Li, D. Gasevic, and G. Chen, “Establishing Trustworthy Artificial Intelligence in Automated Feedback,” Jul. 2022.
- [27] D. Kaur, S. Uslu, K. J. Rittichier, and A. Durrezi, “Trustworthy Artificial Intelligence: A Review,” *ACM Computing Surveys*, vol. 55, no. 2, pp. 39:1–39:38, Jan. 2022. doi: 10.1145/3491209
- [28] D. Gursoy, O. H. Chi, L. Lu, and R. Nunkoo, “Consumers acceptance of artificially intelligent (AI) device use in service delivery,” *International Journal of Information Management*, vol. 49, pp. 157–169, Dec. 2019. doi: 10.1016/j.ijinfomgt.2019.03.008
- [29] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, “User Acceptance of Information Technology: Toward a Unified View,” *MIS Quarterly*, vol. 27, no. 3, pp. 425–478, 2003. doi: 10.2307/30036540 Publisher: Management Information Systems Research Center, University of Minnesota.
- [30] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei *et al.*, “Llama 2: Open Foundation and Fine-Tuned Chat Models,” Jul. 2023, arXiv:2307.09288 [cs].
- [31] H. Luo and L. Specia, “From Understanding to Utilization: A Survey on Explainability for Large Language Models,” Feb. 2024, arXiv:2401.12874 [cs].
- [32] I. O. Gallegos, R. A. Rossi, J. Barrow, M. M. Tanjim, S. Kim, F. Derroncourt *et al.*, “Bias and Fairness in Large Language Models: A Survey,” Mar. 2024, arXiv:2309.00770 [cs].
- [33] N. Dainton, “1 Bedeutung und Wert von Feedback,” in *Feedback in der Hochschullehre*, ser. utb-Titel ohne Reihe. Haupt, Dec. 2020, pp. 11–22. ISBN 978-3-8252-4891-8
- [34] J. Fengler, *Feedback geben: Strategien und Übungen; ... mit über 100 Übungen*, 4th ed., ser. Beltz Weiterbildung. Weinheim Basel: Beltz, 2009. ISBN 978-3-407-36471-5
- [35] R. de Villiers, “7 Principles of highly effective managerial feedback: Theory and practice in managerial development interventions,” *The International Journal of Management Education*, vol. 11, no. 2, pp. 66–74, Jul. 2013. doi: 10.1016/j.ijme.2013.01.002
- [36] J. Ade and U. Gläßer, “Lehrmodul 12: Feedback in der Mediation,” *Zeitschrift für Konfliktmanagement*, vol. 12, no. 2, Jan. 2009. doi: 10.9785/ovs-zkm-2009-60
- [37] N. Dainton, “2 Wo klemmt es,” in *Feedback in der Hochschullehre*, ser. utb-Titel ohne Reihe. Haupt, Dec. 2020, pp. 23–32. ISBN 978-3-8252-4891-8
- [38] A. Seidl, “Dein Wunsch geht in Erfüllung,” in *Freundlich, aber bestimmt – Die richtigen Worte finden in Gesundheitsberufen*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 51–70. ISBN 978-3-642-41803-7 978-3-642-41804-4
- [39] A. Dohrenwend, “Serving Up the Feedback Sandwich,” *Family Practice Management*, vol. 9, no. 10, pp. 43–46, Nov. 2002.
- [40] A. J. Henley and F. D. DiGennaro Reed, “Should You Order the Feedback Sandwich? Efficacy of Feedback Sequence and Timing,” *Journal of Organizational Behavior Management*, vol. 35, no. 3–4, pp. 321–335, Oct. 2015. doi: 10.1080/01608061.2015.1093057
- [41] J. Sammet and J. Wolf, “Präsenztraining im Blended Learning,” in *Vom Trainer zum agilen Lernbegleiter: So funktioniert Lehren und Lernen in digitalen Zeiten*, J. Sammet and J. Wolf, Eds. Berlin, Heidelberg: Springer, 2019, pp. 55–65. ISBN 978-3-662-58510-8
- [42] J. Grenny, K. Patterson, R. McMillan, A. Switzler, and E. Gregory, *Crucial Conversations*, 3rd ed. New York: McGraw Hill, 2022. ISBN 978-1-260-47419-0
- [43] F. Eidenbenz, “Standards in der Online-Beratung,” in *Handbuch Online-Beratung*. Vandenhoeck & Ruprecht GmbH & Co. KG, Nov. 2009, pp. 213–228. ISBN 978-3-525-40154-5

A framework for enabling ex-ante social impact assessment of project-based technological solutions: the case of Remote Infrastructure Inspection

Nikolay Zherdev¹, Olivier Klein², Umberto Sconfienza³, Philippe Gerber⁴, Daniel Vladušić⁵, Jethro Butler⁶, Aljosa Pasic⁷

^{1,2,3,4}Luxembourg Institute of Socio-Economic Research, 11, Porte des Sciences, Esch/Alzette, Luxembourg. ORCID: 0009-0008-8622-7308, 0000-0001-6661-9883, 0000-0002-8312-6631, 0000-0003-2324-5708. Emails: Nikolay.Zherdev@liser.lu, Olivier.Klein@liser.lu, Umberto.Sconfienza@liser.lu, Philippe.Gerber@liser.lu

⁵XLAB d.o.o., Pot za Brdom 100 / SI - 1000 Ljubljana, Slovenia. ORCID: 0000-0001-6032-8575. Email: Daniel.Vladusic@xlab.si

⁶University of Warwick, Coventry CV4 7AL, United Kingdom. ORCID: 0000-0002-9388-0451. Email: Jethro.Butler@warwick.ac.uk

⁷BDS Global Strategy & Innovation, Albarracín 25, Madrid, Spain. ORCID: 0000-0003-0150-5732. Email: Aljosa.Pasic@eviden.com

Abstract— Social Impact Assessment (SIA) is the systematic examination and management of both the intended and unintended social consequences, encompassing positive and negative outcomes, resulting from designed interventions (such as policies, plans, or projects) and any social changes instigated by these interventions. In this paper, we present a strategy to define and validate social impact indicators incorporating participatory approaches into the general impact assessment framework. The paper reports on the first results of an ongoing SIA developed for the evaluation of the impact produced by a Remote Infrastructure Inspection (RII) toolset developed to increase the resilience of critical infrastructures within the framework of the SUNRISE Horizon Europe project. Several stages of the indicators' selection procedure were proposed to ensure the validity of the selection. Our approach is then applied to identify social impact subcategories within the RII Toolset, aimed to introduce less effort-consuming ways of inspecting typically large infrastructures.

Index Terms— Social impact assessment, impact indicators, remote infrastructure inspection.

I. INTRODUCTION

THE COVID-19 pandemic revealed that European societies highly depend on the uninterrupted supply of essential services during global crises, especially during pandemics. These essential services encompass various systems such as energy and water supply, transport networks, health and social services. These services are referred to as critical infrastructures (CIs) due to their importance to society.

The resilience of critical infrastructures (CIs) during such events is crucial for maintaining societal stability, public health, and economic continuity. In this context, resilience refers to the ability of critical infrastructure systems to prepare for, withstand, recover from, and adapt to adverse conditions. It is of paramount importance that European CIs possess resilience, the capacity to adapt to evolving risks and the ability to swiftly recover from both anticipated and unexpected disruptions. However, European CI operators and public authorities are still not fully prepared to address the risks as the threats are swiftly evolving, while the interconnectedness of CIs in Europe has become increasingly complex and digital. In practice, this implies that disruptions in a single CI can increase social, economic, and environmental impacts across other CIs, regardless of the sector of operation.

The SUNRISE project was funded by the EU to address these issues, ensuring the availability and continuity of critical services in Europe through the introduction and deployment of several technological solutions aiming to provide greater adaptability and reliability of CIs. The project, running from October 2022 to September 2025, involves 41 partner organizations, including public and private CI operators, authorities, and technology developers.

Any technical solutions, depending on their purpose and way of operation have the potential to affect the well-being of various groups of people, i.e. individuals, communities, and the society as a whole. At times this impact is direct and intentional: e.g., a technical solution to enhance the economic performance of a company has a direct impact on its financial

¹This work was funded by the European Union through the Horizon Europe research programme under grant agreement n°101073821 (SUNRISE)

indicators. However, other kinds of impact are not directly visible or identifiable: e.g., teleworking tools while allowing performing professional duties from various locations, may affect the communication patterns of the employees their work-life balance, and other social aspects, such as the emergence of a feeling of social isolation.

According to The International Principles for Social Impact Assessment, SIA can be defined as “the processes of analysing, monitoring and managing the intended and unintended social consequences, both positive and negative, of planned interventions (policies, programs, plans, projects) and any social change processes invoked by those interventions” (p.5) [1]. This definition suggests that social impacts encompass a broad range of planned and unplanned consequences caused by a project, while the affected people may include various levels of stakeholders ranging from individuals, and social and professional groups, to the communities and society at large [2]. Besides, SIA is crucial to project organisers, reinforcing their social missions, objectives, and strategies [3].

Thus, the need for a balanced assessment of social aspects of planned intervention is of paramount importance to consider carefully all the possible groups affected and define the key areas of impact through a thorough estimation of indicators, which provide measurable and objective data for informed decision-making. Otherwise, the outcomes of inaccurate evaluation - e.g., priority areas, predicted impacts, measures of significance, etc. - could be either biased or misleading, and may not necessarily reflect stakeholders' values and beliefs [4]. Given that the setting of each project, specifics of proposed solutions and nature of impact vary a lot for each planned intervention, SIA cannot be initiated with a ready-set checklist of possible impacts. It is necessary that the set of impacts is identified “from an awareness of the project and an understanding of how the project will affect what is important to the project’s stakeholders” (p.2) [2].

While an increasing number of projects consider the importance of SIA, many of them encounter difficulties in choosing the most suitable strategy and set of methods that align with the unique requirements of a project [5]. A significant challenge for projects lies in the determination of what impact metrics to apply, and what indicators to measure and report [6]. This complexity extends to the manner in which SIA indicators are defined and verified, especially in the context of planned implementation of technological tools not yet in place. There is a notable research gap in the current approaches, which are often highly case- and technology-specific in relation to indicator lists' definition, paying little attention to the processes of identifying impact categories and validating their results [7]. There is another research gap in the existing SIA methods and frameworks that support technological solutions, as the ex-ante stage of the assessment is often regarded as a subsidiary component of the assessment [8].

While this stage can be regarded as initial and not bringing the final results, it establishes a foundation for further steps

and forms an impact category framework for the whole SIA of a project.

The objective of the study is to identify how an adaptive and participatory SIA framework can be developed to inform and improve the deployment of project-based technological solutions not yet deployed. The article presents the results of the framework elaboration process, discussing the set of assessment stages adaptive to the requirements and limitations of the project and adjusted to the toolset's development stages. These stages enable the implementation of a specific yet versatile set of consecutive and analytical steps. Using the example of one technological solution developed within the SUNRISE project, the paper discusses the first stage of assessment, which allows the provision of feedback to key stakeholders and adjustment of the tool prior to the testing and deployment stages. The ex-ante evaluation procedure advocates the assessment based on the project scope, outputs, and stakeholders directly involved in tool development, defining and verifying the relevance of impact areas and dimensions. Building upon these focal points, interim results of the assessment are discussed with a set of key issues identified to increase the social sustainability of the project.

The article is organized in the following way. Section 2 provides background information on the SUNRISE project, presents a specific technological tool of this research (a tool for RII) and reviews relevant literature on social impact assessment methodologies and theories. Section 3 sets out the general SIA framework of SUNRISE, explains the ex-ante strategy used to conduct the social impact assessment, and discusses data collection methods (surveys and focus groups) applied. Section 4 presents the findings of the SIA and describes the identified social impact categories of the project. Section 5 interprets the results of the ex-ante assessment in the context of the project and their social implications, provides an analysis of the significance of the identified social impacts, discusses the potential implications for stakeholders and affected communities, and addresses limitations and challenges encountered during the assessment. Section 6 concludes by summarising the main findings of the SIA, suggesting practical applications, and areas for further investigation.

II. THE CONTEXT OF THE STUDY

A. The SUNRISE project

In the context of COVID-19, European CI operators and public authorities have been faced with numerous challenges in managing the risks of future crises, as threats continue to evolve rapidly [9]. These challenges include the lack of collaboration among CIs; the absence of thorough risk evaluations; a scarcity of specific strategies and measures to maintain operations during a pandemic; and a shortage of resilience-oriented tools available before and during the crises. To address these challenges, the SUNRISE project has targeted five key objectives¹:

¹ <https://sunrise-europe.eu/about>

- To promote dynamic cooperation among European CIs, spanning various sectors, and involving both public and private entities.

- To identify services and CIs that are crucial during a pandemic, understand their interconnections and dependencies (including risks of cascading effects), and devise effective mitigation strategies.

- To produce a strategy and a suite of advanced technologies to ensure CI resilience and business continuity during a pandemic.

- To test the newly developed strategy and technologies in real-world settings throughout Europe.

- To promote a resilience-oriented approach within and across European borders.

The project's consortium includes 18 public and private CI operators and authorities from various sectors (energy, transport, health, digital infrastructure, water supply, and public authorities) and countries, including EU Member States as well as EU-associated countries.

The project aims to develop a suite of new technologies and applicable solutions, including Risk-Based Access Control tool (a tool that minimizes risk when accessing critical infrastructure in a scalable, privacy-preserving manner; Resource Demand Prediction and Management tool (a flexible tool designed to handle changing demands for resources during emergencies, regardless of the specific critical infrastructure); Cyber-Physical Resilience tool (this tool detects anomalies, issues alerts for incidents, provides appropriate responses, and conducts real-time risk assessments for critical infrastructure); and Remote Physical Infrastructure Inspection tool (a tool that uses satellite images, unmanned aerial vehicles – UAVs – with various sensors, and performs machine learning identification to detect anomalies and thus to continuously inspect physical infrastructure)².

B. A tool for Remote Infrastructure Inspection

The latter technological toolset (RII) was selected as a case study for the present article for several reasons. First, it combines the modules (discussed below) that are increasingly raising concerns about their application in relation to social impact [10] and potential misuse of obtained information [11]. Second, the societal advantages of obtaining high-definition data should be thoroughly balanced against societal concerns such as privacy [12], which implies an overarching prominence of an appropriate SIA for such technological solutions. Moreover, the ethical use of RII solutions is highly dependent on the setting in which they are used [13], which highlights the prominence of impact evaluation and management, specifically within the context of an international project and multiple CIs involved (seven utilities in transportation, water and energy supply, and telecommunication) in the RII tool's development, testing, and deployment. Given that the SIA framework proposed in this article was created for the entire suite of SUNRISE tools (four technological solutions), its application in such a specific context of RII informs the

manner in which it might be adapted to other settings, demonstrating the replicability on other technological solutions.

RII tool, developed in the context of SUNRISE has as a main aim to introduce less effort-consuming ways of inspecting typically large infrastructures (which necessitates the allocation of considerable resources to the physical infrastructure inspection). To do so, the tool aims to provide a more comprehensive view of the infrastructure and its surroundings, allowing inspectors to assess large areas efficiently and frequently. It consists of two main approaches, which are divided according to their vicinity to the infrastructures: the use of satellite imagery for infrastructure inspection (from afar) and the use of UAV (drone) imager to inspect the infrastructure. Both modalities are supported by artificial intelligence (AI) tools, able to sift through the acquired data and provide the focus on potential anomalies, thus alleviating the CIs of manually checking the acquired data. Furthermore, the use of certain modalities (e.g. satellite imagery) with the pre-trained AI models allows for efficient prediction of certain features (e.g. vegetation height, which needs to be controlled to comply with clearance requirements and ensure safety conditions), that are very hard to predict manually. In this way, the infrastructure is continuously monitored in a non-invasive manner.

Typically, the entire infrastructure is controlled by satellite imagery, responsible for detecting any adverse events. However, satellite imagery, despite very high spatial resolutions, remains limited in the precision of results it can provide for the inspection of critical infrastructure. Thus, they are used as a trigger to activate the more costly and labour-intensive UAV-based inspection, which offers a more targeted, localized area-focused monitoring of regions of interest. UAV-based methods work on the acquired imagery and offer a more detailed view of the state of the infrastructure. The AI involved in both methods is from the field of Computer Vision and is capable of predicting the height of vegetation, and detecting changes and states, including the condition of insulators in critical infrastructure. In this sense, it points out anomalies, avoiding the need for human operators to carry out continuous and tedious inspections in the field, and enabling them to focus directly on interventions. As a result, this AI-based tool optimises inspections and interventions.

By introducing these tools, various social aspects are expected to be impacted (e.g. safety, well-being, and privacy), highlighting the importance of conducting a comprehensive SIA.

C. Review of relevant literature on social impact assessment theories and methodologies

SIA is a discipline within social science that aims to provide insights into the vulnerabilities, risks, capacities, well-being, and resilience of various groups of people. It offers valuable information for designing planned interventions that enhance potential risk reduction, resilience, and social sustainability [14] – [17].

² <https://sunrise-europe.eu/about>

SIA involves the evaluation, monitoring, and management of both positive and negative social consequences that projects may have on the well-being of individuals, specific groups of people (e.g. workers) and communities [2]. The development of the SIA discipline has been closely associated with a project-oriented context, applied to enhance positive outcomes and minimize negative social effects of planned interventions [18]. To provide an appropriate evaluation, the SIA process is recommended at each key project phase [2]. Therefore, impact assessment practitioners should consider the specific context of a project-based evaluation (specificity of a project form, e.g. phases, timeframe, location, etc.) as well as intervention characteristics (product, technological solution, physical construction, etc.). The following paragraphs present the SIA specificities of a project and remote infrastructure technologies to which the SIA will be applied.

A project is recognized as a form of transient organisational structure that is becoming increasingly prevalent in contemporary society, where “resources are assigned to undertake a unique, novel and transient endeavour managing the inherent uncertainty and need for integration in order to deliver beneficial objectives of change” (p. 7) [19]. A project, during its lifespan, adheres to a set of meanings, holds specific values, and operates on a philosophy shaped by the consortium's relationships [20], usually reflected in the project's guidelines and proposals. According to Bakker et al [21], the success of a project fully depends on the precise definition of priorities, management procedures, and strategic methods to address issues effectively. Such effectiveness should be understood in broad terms, considering changes beyond overcoming the technical challenges (reaching defined goals), and encompass various areas that the planned intervention may affect, including the social domain.

Many studies have discussed the evaluation of the social impact within projects [22]. According to Passani [23], within the context of the organisational form of a project, an appropriate impact assessment strategy (including SIA) requires analytical rigour to ensure the validity of the evaluation; several iterations of assessment throughout all phases of the project; appropriate data collection methodology accommodating the analysis of differences in impacts; and specific budget to ensure proper execution of the assessment.

In order to consider these general requirements, a proper project-based SIA framework needs to be set up. According to Vanclay et al [2], the established good practice SIA comprises four phases: understanding the issues; prediction, analysis, and assessment of the likely impact pathways; development and implementation of strategies; and design and implementation of monitoring programmes. The authors highlight that due to the context of a project, these phases are “somewhat sequential, but which also overlap” (p. 7) and need to be adjusted (ibid). Indeed, as each project has its own context and setting it is hard to present a pre-defined framework that fits any planned intervention, the modifications in more precise stages of assessment as well as methods selection need to be made by SIA practitioners. In particular, corrections of an

ongoing evaluation deserve special attention: “Through data collection and analysis, SIA is a learning process, and consequently initial assumptions and preliminary understandings may need to be modified in the light of new information, so there needs to be an iterative process of validation and update informed by an on-going process of consultation with project proponents and other stakeholders” (ibid, p.7). Such a perspective is relevant for all the dimensions of a project (e.g. its scope, area, stakeholders, etc), however in the context of SIA framework elaboration, where execution steps should be based on a realistic understanding of the project's potential impacts [23], impact areas identification and validation take on particular significance.

D. The social impacts of Remote Infrastructure Inspection solutions

While the technical effectiveness of UAVs and satellite imagery to tackle operational issues of CIs is rather evident, this alone should not lead to their implementation regardless of the social impact they may cause. According to Pastra et al [24], such technological solutions may induce both positive and negative effects on people, while the corresponding analysis should go beyond physical impact (e.g. collision risk) and be extended to include other types and forms of possible effects. However, the existing studies on the social impact of RII-related tools are highly limited, focusing predominantly on the social effects of UAV usage (in contrast to satellite imagery). In relation to the latter aspect, Sandbrook [25] distinguishes four main categories of potential social impacts in relation to drones' usage (studying the case of biodiversity conservation): safety, privacy, psychological well-being, and data security. The first category refers to potential safety benefits to those individuals on the ground in the event of a UAV crash in comparison to larger, manned aircraft [26], the presence of safety mechanisms, and the potential to detect criminal activity [25]. At the same time, Lee et al [27] stress that the absence of a pilot can make UAVs more prone to crashes. From the privacy perspective, the application of drones (as well as other aerial monitoring and image-capturing technologies, including high-resolution satellite imagery) raises various ethical concerns in relation to possible surveillance issues, violation of privacy, human rights, and civil liberties [28]. The psychological well-being of individuals might be enhanced due to empowering local communities by providing access to their own data-gathering tools [29], while UAV applications may provoke fear and confusion among those on the ground [30]. Finally, the usage of drones may raise concerns about how collected data is used, to what extent is it protected and to whom is it shared [25]. The latter aspect is specifically important considering the implementation of AI components in the RII tools (which raises ethical and data privacy considerations) [31].

However, these broad and relevant categories do not consider other possible impacts and specific social groups related to differing purposes of UAV usage, as in the case of CIs setting. This limitation implies the creation of a context-specific

set of categories tailored to the RII tool proposed by the SUNRISE project.

III. METHODS

The framework presented in this section aims to evaluate the impact of the project's results and form a sustainability roadmap for their continued existence beyond the project's end. The project requirements imply the identification of the societal impacts of SUNRISE solutions and communication of the corresponding set of recommendations that will help the solutions' adopters (CI operators) to mitigate potential negative impacts and enhance the benefits of the SUNRISE tools, given that foreseen impacts will be evaluated continuously during the project. The section begins with an overview of the overall strategy applied for all the technical solutions of SUNRISE, followed by the ex-ante assessment methodological framework applied, providing an overview of a more detailed explanation of methods used in relation to the social domain in relation to the RII tool.

A. Framework

This section presents a comprehensive strategy for continuous evaluation of the impact produced by SUNRISE toolsets. This approach advocates initiating the evaluation by defining the project scope, outputs, stakeholders involved, and relevance of various impact areas and dimensions. Following this, an impact evaluation process is outlined detailing its strategy, phases, timing, and limitations. Subsequently, the data-gathering process starts, utilizing the methods selected in the previous phase. Once the data is gathered, the analysis ensues and provides a comprehensive report on the findings. Moreover, given the necessity for continuous impact evaluation to ensure the quality of SIA, two additional stages are included in the strategy assessment during tools' testing and

monitoring, followed by final reporting. The strategy aims to provide a specific yet versatile set of consecutive and analytical steps, which includes the stages presented in Table I.

To improve the quality of the assessment, and to incorporate the most applicable assessment practice that focuses on participatory approaches, the present strategy prioritises the project's specifics, considering the solutions characteristics proposed by SUNRISE and their development stages. In this article, the RII tool is presented in detail, presenting the methodology of defining and validating indicator categories for the tool. As the tool is not yet in a testing phase within the CIs, the aim at this stage of the project is to present the progress of the evaluation strategy, represented by the results of identification of the ex-ante impacts and validation of the results by confronting the opinions of the developers and those of the CI representatives.

The following section presents a detailed explanation of the methodological framework applied for social impact evaluation. As the first two stages are highly project-specific and out of the scope of the study, the article focuses on defining and validation of indicators, corresponding to stages 3 and 4 of the General Framework presented in Table I.

B. Defining and validating the SIA indicators (Stage 3)

The key issue of impact evaluation lies in defining the appropriateness of the indicators, such that the coverage of all impacted domains and parameters is ensured. To do so, the authors applied an approach drawing upon the '3S' methodology (self-validation, scientific validation, and social validation) [32], which emphasizes reliance on expert judgments and public participation. However, the adapted methodology has been refined to facilitate a suitable ex-ante assessment, recognizing the critical role of expert knowledge in the early stages of solution development. Public

TABLE I.
SIA'S GENERAL STRATEGY FOR THE SUNRISE PROJECT

Stage	Description
1. Screening and Profiling	Understanding potential issues, gathering secondary data on impact areas, and engaging experts and partners (help to formulate an overall evaluation strategy, establishing a common framework that can be applied to other impact analyses).
2. Scoping	Characterising the technological solution, involving experts and developers (identification of limitations and areas requiring more in-depth evaluation based on participative approaches to create a list of evaluation indicators tailored for each area).
3. Defining and validating impacts	Defining (and selecting) specific indicators and variables within the social area.
4. Verification of indicators and Assessing Foreseen Impacts	Investigating foreseen impacts across the evaluation domains, outlining expected changes ex-ante as perceived by various stakeholders in relation to identified indicators.
5. Interim reporting and Provision of feedback to allow Enhancement and Redesign	Providing feedback to stakeholders about potential benefits and issues related to Tools' deployment and usage (support to improve the toolsets from a social perspective).
6. Assessing impacts and cumulative effects	Investigating direct and associated impacts identified during tools' testing (with the use of both qualitative and quantitative approaches).
7. Monitoring and final reporting	Ensuring the accuracy of the assessment and including the development of the sustainability roadmap.

participation is integrated into later stages of the SIA, once the tool is deployed and tested by additional groups of end-users. Furthermore, expert validation is integrated with ex-ante impact evaluation to streamline the process and ensure more efficient allocation of project resources. This approach minimizes unnecessary data collection steps, crucial in project-based impact assessment contexts where resources are often limited.

The approach is based on two consecutive stages. First, the initial validation phase of defining the indicator categories uses information obtained from a literature review to create an initial set of indicators, which forms a basis for their later refinement. The second validation phase is performed by the assessment team itself to capture social sciences and humanities (SSH) partners' opinions on the set of possible parameters and to refine the initial list of indicators. It is worth noting that there is a necessity to balance the relevance of potential impacts and the inclusivity of indicators.

To perform the identification stage, the authors applied the indicator-based approach presented by Schuck-Zöllner et al [33]. This approach implies an 'evaluation cascade', proposing a hierarchical categorisation of indicators, ranging from generalised groups to more precise evaluation units and parameters, containing the following categories: Dimension – Criterion – Indicator – Method. Combined, these categories form a classification framework for the social area.

C. Verifying the SIA indicators and conducting an ex-ante assessment (Stage 4)

Once the initial validation (defining) of indicators is finalised, the expert participant verification stage is applied to reach a consensus on the key impact parameters as well as to ensure transparency in the assessment procedure. This validation stage is needed to ensure the credibility of SIA and to avoid false assumptions that might occur during the initial validation process.

To streamline the evaluation procedure, this phase is combined with data gathering for the assessment process. This stage comprises two consecutive steps: (i) conducting an expert-oriented survey and (ii) carrying out a focus group to provide a better understanding of key issues identified by the survey. The expert stakeholder groups involved in the evaluation process were formed independently for each tool, while the participants were selected based on their expertise in the technical and operational specifications of technological solutions. These expert groups are comprised of tool developers and CI representatives who are directly involved in the tool development and implementation processes and therefore they possess expert knowledge of the solution's components, deployment and operation procedures, which is necessary to conduct an ex-ante evaluation. Finally, the framework's timespan and methods are tailored to align with the toolsets' development phases.

(i) Preparation and conducting a mixed-method survey on foreseen (ex-ante) impact assessment for the tool: this phase is needed to further advance indicators' validation and to identify key impact parameters (by juxtaposing response

trends from two aforementioned groups of stakeholders involved in the tool's deployment).

(ii) Carrying out a participatory online focus group for the tool (involving CI representatives, and tool developers) is required to assess the foreseen impact on key Impact Assessment topics based on the survey results. Focus groups are used to gather additional qualitative data by engaging a small group of participants in an open discussion about specific topics. These tools are particularly useful for exploring attitudes, opinions, and perceptions. Besides, focus groups are helpful to generate a range of perspectives, and deeper insights, reveal nuances, and capture a variety of viewpoints by stimulating discussion among participants [34]. SIA practitioners may choose to conduct surveys and focus groups sequentially, using the findings from one phase to inform the design or focus of the subsequent phase. The recordings of the moderated online focus group were transcribed for subsequent analysis of the impact assessment categories.

These steps of defining and verifying the indicators are presented in Figure 1. Carrying out these steps was followed by




Fig 1. The framework for defining and verifying the SIA indicators for the SUNRISE technical solutions

formulation of feedback to consortium partners to mitigate negative effects and enhance positive impacts.

IV. RESULTS

The following parts present categorised sets of indicators to assess the social impacts of the tools. These sets represent

the results of the two phases of the indicator validation framework introduced in the previous section.

A. Defining the indicator categories for the Remote Physical Infrastructure Inspection tool

To initially validate indicators, we utilized an inclusive set of social impact assessment dimensions and criteria based on Rainock et al. [35], adapted into 'Dimensions' and 'Criteria' categories following the framework by Schuck-Zöller et al. [33]. The original set included seven dimensions: Stratification, Employment, Health and well-being, Human

rights, Networks and Communication, Experience of conflict and crime, and Cultural Identity and Heritage. During the second phase of indicator verification, it was found that only a limited number of these dimensions directly applied to the RII technical solution. Consequently, the categories of Cultural Identity and Heritage, and Experience of conflict and crime were deemed out of scope to streamline the assessment process. Additionally, the dimension of 'Stratification' was refined to focus specifically on equality-related impacts relevant to the project's scope.

TABLE II.
DIMENSIONS, CRITERIA, AND INDICATORS USED FOR THE SIA OF THE REMOTE INFRASTRUCTURE INSPECTION TOOL

Toolset description		Ethics, Human Rights and Privacy		Equality	
Toolset deployment details	Type of End-Users, Type of Disasters mitigated, Type of CI sectors, Countries of planned deployment, Existing Measures	Regulation compliance	GDPR compliance	Social equality	Gender, Workgroups, Physical and Health conditions, Spatial inequalities, Temporal inequalities
Toolset outcomes	Lockdowns: Effects on Lockdowns Number, Effects on Lockdowns Length Access to services: Quality and availability of services, Service continuity, Equitable distribution of services, Affordability of access to services, Service barriers, Service Delivery models, Technology adoption	End-User involvement Data management	End-users' participation Personal data processing, Special categories of data, Intrusive methods of Data gathering, Risk of identification of individuals, Non-public data usage, Data Export to non-EU countries, Risk to rights and freedoms, Incidental Findings	Skills	Skills and training, Training requirements and amount
Employment		Health and Well-being		Decision-Making, Networks and Communication	
Work environment	Change in Work Environment, Absenteeism rate	Safety and Security	Safety incidents, Security incidents, Emergency response, Long-Term dynamic, Risk-Mitigation Strategies, Technological Testing and Validation Evaluation, False alarms, Recovery Time and Cost, Correction Actions, Public Health Interventions, Infrastructure Security, Emergency planning	Communication practices	Information flows and information dissemination, User Awareness regarding features, Accessibility and Inclusivity, Scope and extent, Frequency and timeliness, Two-way communication, Collaboration and Knowledge Sharing, Usability and User Experience, Prediction information sharing within the CI, Prediction information sharing with stakeholders, Transparency with stakeholders, Information sharing with customers, User Engagement, Inspection information sharing within the CI, Inspection information sharing with stakeholders
Work-life balance	Change in Work-Life Balance, Autonomy, Flexibility, Working remotely	Health (Mental health, Physical health, Health risks)	<u>Mental health</u> : Stress and Anxiety, Social Isolation and Loneliness, Sleep Quality, Emotional Well-being at workplace, Cognitive Load and focus, Tool Dependency, End-user behaviour <u>Physical health</u> : Physical Activities, Sedentary behaviour, Infection prevention at workplace Health risks: Adverse Health Effects, Injury rates, Emergency response mechanisms, Exposure to Health risks for end-users, Healthcare goods distribution, Access to Healthcare services, Health risks for CI customers	Interpersonal relations	Frequency, Relationship strength, Community building, Networking opportunities, Inclusivity, Belonging, Conflicts
Working hours	Change in Working Hours, Working time, Workload, Working outside office hours, Efficiency at work			Decision-making and Management	User Empowerment, Distribution of decision-making authority, Prioritisation of user/customer needs, Collaboration and co-creation, Delegation of management tasks, Decision-making transparency, User-initiated actions, Distribution of management responsibilities, Management during peak demand periods, Policy development

The outcomes of the two initial validation stages are presented in Table II, listing indicators linked to dimensions and impact criteria for assessing the RII tool's social performance. It includes specific dimensions and corresponding Criteria that evaluate potential positive and negative impacts on diverse social groups. Through the indicator-definition process, five key dimensions of social impact were identified: 'Ethics, Human Rights and Privacy', 'Equality', 'Employment', 'Health and Well-Being' and 'Decision-Making, Networks, and Communication'. Additionally, a general 'Tool Description' dimension was necessary to account for the toolset parameters and deployment factors that can significantly affect the social domain.

B. Verifying the SIA indicators for the Remote Infrastructure Inspection tool

The analysis of ex-ante impacts in relation to the implementation and usage of the RII tool aims to provide an evaluation of the potential future effects of the deployment and usage of the tool. In the first phase, eight survey contributions were collected, including five respondents involved in the tool development and three representatives of CIs (transport sector, water and energy supply sectors). In the second stage, the focus group was conducted, involving twelve participants, including representatives from five CIs concerned with the RII tool deployment (transport sector, telecommunication sector, water and energy supply sectors). The focus group structure was aligned with the 'Dimension' parameter of the list of evaluation indicators presented in Table II, covering all the dimensions identified in the 'Defining and verifying' stage.

The following sub-sections present four key impact categories identified during the validation and verification procedures of SIA, namely awareness of pre-existing infrastructure inspection solutions, data collection and management, work satisfaction, and safety.

Awareness of pre-existing infrastructure inspection solutions

Tool developers had difficulty answering whether there were pre-existing measures prior to tool application, while only CI operators were able to confirm the existence of prior solutions. Focus group participants from tool development partners confirmed their limited awareness of the pre-existing solutions. At the same time, a focus group member from a CI highlighted that not all the developers are aware of the existing measures in CIs.

Data collection and management

The operational functioning of the inspection toolset involves potentially intrusive methods of data collection and processing, including various satellite imagery modalities (optical, multispectral, etc.) and the use of UAVs equipped with different visual sensors (optical, multispectral, etc.). According to survey contributions, there will be no intentional recording of people. However, since certain infrastructures in areas frequented by citizens may be recorded, it is the responsibility of tool operators to pixelate any potentially identifying elements in the footage. Thus, according to survey results, there will be no risk of identifying individuals. Regarding

anonymization practices, Tool developers confirmed that they are only needed for UAVs (in contrast to satellites, where such practices are not needed), and they might be advanced beyond face pixelization. Regarding the anonymization procedure, a CI representative described it as a straightforward task, but clarification of the procedure is needed. Another concern expressed during the focus group refers to the decrease of the toolset's functionality due to the excessive anonymisation techniques applied (as in case of over-blurring), which may increase false positives by blocking significant parts of the images provided by UAVs. According to a CI representative, a potential solution to this issue might be represented by the application of the tool model prior to the anonymisation of the images.

Work environment

Although most survey respondents foresee no changes in the Work Environment, three participants noted potential impacts. They anticipate improved employee satisfaction due to safer remote work conditions and fewer on-site inspections. This could also reduce absenteeism by around fifteen to twenty per cent, as discussed in focus groups, primarily attributed to enhanced safety and risk minimisation."

Additionally, changes in employees' working hours may occur due to seasonal adjustments and reduced work hours with maintained financial compensation. Respondents anticipate reduced working hours and lighter workloads, alongside increased efficiency. This trend was supported in focus groups, emphasizing improved efficiency and inspection processes through detailed infrastructure information, minimizing unnecessary on-site visits and equipment requirements. Current inspection practices often involve regular visits regardless of immediate issues, whereas future approaches may prioritize visits only during operational failures.

Survey respondents indicate a mixed impact on employee job satisfaction. Positive changes include Increased Engagement, Improved Work-Life Balance, Enhanced Recognition, Career Advancement Opportunities, Meaningful Job Responsibilities, Positive Team Dynamics, Supportive Leadership, and Job Security. Conversely, potential negative impacts noted are Unrealistic Expectations, Increased Workload, Lack of Recognition, and Ineffective Leadership. Focus group discussions confirmed these trends, highlighting both positive and negative aspects. For instance, there were concerns about unrealistic expectations related to satellite image capabilities and optimism regarding improved recognition and employee well-being.

Safety and security

According to the major trend of the verification stage responses, the contributors foresee a positive impact on the safety and security of the end users, which includes decreasing in the number of workspace safety and infrastructure security incidents, as well as an increase in emergency response efficiency. At the same time, the verification stage is indicative of the positive foreseen impact that can be identified in the number of false alarms. Besides, the development team representatives marked the high level of dependency of the number of false alarms on the CIs'

preferences, as some of the infrastructures prioritise the detection of any kind of changes, while the other CIs are interested in the detection of only a certain type of changes.

The focus group identified differing opinions regarding the number of security incident detections. CI representatives suggested there might be a decrease in the number of detections due to earlier problem identification, whereas tool developers argued that the number of incident detections might increase due to more frequent checks that would lead to a reduction in the number of security incidents. However, all the participants unanimously agreed that a positive trend in infrastructure security will prevail.

Concerning the efficiency of the Emergency response procedures, the focus group participants highlight the increase in speed of feedback and reaction time, as the tool implies real-time updates of the required information.

V. DISCUSSION

This section reviews the results of the two initial ex-ante evaluation procedures outlined in this paper, namely defining and validating indicators (Stage 3 of Table I) and verifying indicators' categories (Stage 4 of Table I). These procedures establish the initial framework stage of impact areas' definition and evaluation, forming a basis for the following stages of continuous SIA of the project.

A. Defining and validating the indicators (Stage 3)

In the stage of defining categories and specific impacts, six inclusive impact dimensions were identified: Toolset description; Ethics, Human rights and Privacy; Equality; Employment; Health and Wellbeing; and Decision-making, Networks and Communication. While some of these dimensions are expected for SIA of the RII solutions (e.g. ethics concerns, health-related consequences), other dimensions of possible impacts are not commonly present in SIA literature of the RII technological tools (e.g. effects on decision-making and employment) [25-28]. The assessment of the impact dimensions identified is not consecutive, however, the inclusion of the 'Toolset description' dimension as the initial (introductory) step is deemed necessary as it informs the tool deployment specificities and possible outcomes of its implementation. Accordingly, this dimension allows contextualizing the tool deployment setting and the definition of its end-user groups, which facilitates subsequent impact evaluation.

As it follows from the literature devoted to SIA of the RII-related technologies, the dimension of ethics, human rights and privacy raises numerous concerns [25]. The present methodology facilitates the structuring of this dimension in relation to regulation compliance, data management practices and end-user involvement in the tool development and deployment processes, as these criteria were identified as both relevant and inclusive for the context of the RII tool.

According to existing impact assessment literature, the dimension of equality is essential to the SIA of technological solutions [36], however, given a broad range of possible social groups affected by a technological solution, the list of specific categories, containing related indicators is highly context-dependent and needs to be verified carefully. The

defining and validation stage of the proposed methodology identified the relevance of the 'social equality' and 'skills equality' categories for the RII solution of SUNRISE, allowing us to consider the possible impact on equality for different employees groups and CIs customers.

As CI employees are perceived as the primary end users of the toolset, the employment dimension is crucial for the current SIA. Existing literature emphasizes its significance [37,38], but empirical evidence on the impacts of technological tools on employment is limited [39]. In the RII context, our methodology identified Work environment, Work-life balance and Working hours as key criteria for assessment. The adoption of new technology in daily operations directly affects how people work, including the types and volume of tasks. Our ex-ante SIA findings underscore the increased importance of employees' efficiency for CIs representatives (due to technology implementation), potentially leading to greater job complexity [40]. Moreover, the introduction of new technologies may promote a skill-biased shift within organisations, benefiting skilled employees [41]. Consequently, this shift may raise expectations for higher performance among employees, as tasks supported by new tools could be perceived as less time- and effort-demanding compared to the traditional work settings.

The dimension of Health and well-being is covered in the SIA of technological solutions [42]. However, the list of possible health-related impacts and thus indicator categories to include in the assessment significantly depends on the social groups affected by the implementation of a specific tool [43]. Accordingly, the initial selection of categories for the current impact evaluation was preceded by the identification of social groups affected and end-users of the RII solution. The methodology allowed the assessment team to identify Safety and security, and Health criteria (including sub-criteria of mental, physical, and health risks) as the relevant categories. Such an approach considers a wide range of impacts on individual, group and organisational levels.

Networks and communication are often present in the SIA frameworks, however, the usual focal point of the studies concerning these aspects is communication technologies [35]. At the same time, the impact of technological solutions on decision-making is well-studied in management-related literature [44]. Considering the importance of CI employees as the social group directly affected by the RII tool deployment (as well as the potential change in the everyday work scenarios), the dimension of decision-making, networks and communication is deemed necessary for inclusion in the list of indicators of SIA of the RII tool (Table II). Indeed, communication and networking between employees might be directly or indirectly affected by the implementation of new instruments as they change the work environment, e.g. from on-site to remote inspection. The impact-defining stage results indicate that new ways of obtaining inspection data may have repercussions for the division of responsibilities between employees (or groups of employees) and the decision-making process within a CI.

This inclusive set of dimensions (containing related criteria and specific indicators) forms a basis for the following

verification procedure, aiming to identify and evaluate the key impact categories concerning the RII solution implementation.

B. Verification of indicators and conducting Ex-Ante Assessment (Stage 4)

The impact category of Awareness of pre-existing infrastructure inspection solutions refers to the knowledge of specific CI settings among tool developers. The results of the impact verification stage highlight the need for increasing awareness regarding this parameter within the development team. The lack of such awareness may lead to inconsistencies in the development and deployment phases, which in turn may instigate negative outcomes for CI organizations, their employees, and customers. The high level of importance of contextual and procedural information for the consortium partners conditions the inclusion of the 'Tool description' dimension in SIA, especially in a project-based context, where tool deployment settings may vary significantly. Otherwise, highly prominent peculiarities of tool development and implementation might be omitted. Therefore, considering this aspect in the SIA framework enables the provision of valuable feedback to tool developers.

Another impact category that was identified as specifically important for the RII tool refers to data collection and management. Indeed, the wider dimension of ethics and privacy is highly relevant to the social impacts of technological solutions [45] as well as RII-related tools [25]. However, it is necessary to define what types of ethics- and privacy-related impacts are the most relevant for specific use cases and modules (components) of technological solutions. Arguably, one of the possible solutions to this task is the engagement of groups of experts who are actively involved in tool development and implementation as they possess knowledge concerning the specific context of tool deployment. In the case of RII of SUNRISE, these groups are represented by tool development partners and CIs representatives, as the proposed solution contains multiple components and tool implementation settings vary by infrastructure sector and location. In this regard, the results of impact verification indicated that ethics and privacy-related concerns were specifically relevant for UAV usage in comparison to satellite imagery due to the technical specificity of each module. Among those concerns 'Data collection and management' criterion was identified as the key impact parameter of the dimension of Ethics, Human rights and Privacy. The verification process during the focus group highlighted possible issues and ambiguities concerning data collection and management, namely potentially intrusive methods of data collection and anonymisation procedures. Besides, direct communication between tool developers and CIs representatives during the focus group discussion allowed clarification of data collection and management practices and formulation of possible solutions to maintain tool functionality while applying privacy-related procedures.

The criterion of work environment is rather well studied in SIA literature on technological solutions [46]. However, in the case of the RII tools the existing research on this criterion is limited, thus it may be easily omitted in case the impact

category list is based solely on practitioners' perspectives. This claim bears even higher prominence for the impact categories that constitute this criterion, as technical tools for remote inspection differ significantly, and hence their planned and unplanned impacts on the work environment are also highly varying. In the context of the RII tool of SUNRISE, the stage of indicator verification allowed the impact assessment team to highlight two key dimensions of possible impacts: absenteeism rate, and workload changes. Besides, a possible connection was identified between these aspects and the job satisfaction rate of the CI employees. This link is also confirmed by academic literature, as the change of workload to less physically demanding tasks may result in an increase in the satisfaction level among workers [47]. At the same time, a higher level of employee satisfaction with their jobs and work environment may lead to a decrease in the level of absenteeism [48].

Finally, the indicator verification stage indicated the relevance of the 'Safety and security' criterion for both developers and CIs representatives. As follows from the literature review, the notion of safety is rather well-considered in relation to UAV usage [25],[30]. However, the present impact verification methodology highlighted another attribute of this parameter: in contrast to the mentioned research, the focus of the application of drones is compared to the pre-existing setting of inspection, but not to other ways of conducting aerial vehicles or safety in case of a crash. This allowed the identification of potential context-specific safety implications, relevant to the implementation of RII in CIs, in particular, the change of employee's work tasks from on-site to remote inspection is expected to increase the level of safety. Besides, impacts on emergency response and number of false alarms were identified as highly relevant for focus group participants. Arguably, these security-related impacts are deemed prominent due to the specific context of CIs as the tool deployment setting.

C. Ex-ante SIA framework overview

The SIA frameworks are increasingly valuable for project practitioners responsible for the evaluation of a spectrum of social phenomena impacted by technological solutions. However, existing frameworks provide rather narrow empirical evidence for procedures of defining and verifying the categories of social impacts [34], where the results are frequently grounded in the practitioners' or authors' perceptions [1]. As a result, the frameworks outlining several categories of impacts are characterized by a significant level of inconsistency between such impact categories (ibid).

To mitigate such inconsistency, the authors propose focusing on the methodological phases of ex-ante assessment for defining and verifying indicators, rather than the formulation of a ready-to-apply list of indicator categories. This approach ensures that both the tool specificity and the broader context of its development and implementation are taken into account.

The proposed methodological framework allows the collection of the expert perspectives (e.g. CI representatives, tool developers, SSH experts) on the social impacts of the specific tool, as well as identifying the key impact areas, thus

avoiding grounding the impact assessment criteria solely on the impact evaluation practitioners' beliefs and perceptions. This is done through collection and structuring of the broad and inclusive range of social impacts at the initial stage of indicator definition and validation, allowing to contextualise the list of potential impacts both in terms of the project (e.g. end-users, tool deployment) and specific technological solutions adapted to contain an inclusive list of relevant impact dimensions, criteria and indicators. Subsequently, this range of indicators is verified by conducting successively an expert survey and a focus group, highlighting the key areas where the most prominent social impacts are expected.

Such an approach is tailored to consider the technical specifications of targeted solutions, resulting in enhanced relevance and granularity of the SIA procedure and results. This enables comprehensive impact analysis, where a high level of method adaptability allows for the inclusion of specific indicators that are most relevant to the technology, required to enable more meaningful impact assessments [2]. Besides, expert stakeholder engagement in the ex-ante evaluation allows for specific risk identification and targeted mitigation measures before solution deployment.

Concerning the project-based context of SIA, the proposed methodology proposes a solution to enhance resource allocation during the SIA process by combining methods for indicators' defining and verifying stages, which enables cost savings while ensuring the identification and analysis of critical aspects. The adjustment of resource allocation is often prompted by the context of a project, where impact assessment needs to consider tool development and deployment phases as well as material resources' limitations [23].

At the same time, several limitations of such an approach need to be highlighted. First, the current research presents the application of the adaptable methodology tailored to a specific technological solution, thus if applied to other tools it requires to be adjusted in accordance with other projects' aims and technological specifications to ensure the relevance of the SIA results. Similarly, if applied in a different setting, the approach needs to be adapted to specific project requirements (e.g. development stages, timeframes, resources, etc.) and deployment context (e.g. location, stakeholders, etc.). At the same time, modifying the methodology for different projects or technologies can be complex, resource- and time-consuming. The latter aspect includes the requirement of an extensive analysis of relevant theoretical sources and empirical evidence to establish an inclusive yet relevant list of impact assessment indicators as well as a possible requirement to conduct specialized training to apply the tailored methodology, which may incur additional costs. Moreover, as technologies evolve rapidly (and tend to include additional components) [49], a tailored methodology may quickly become outdated, requiring frequent updates to remain relevant (e.g. to include additional impact categories). Therefore, more empirical evidence of the application of the proposed ex-ante approach to other tools and settings is needed to inform the level of applicability and transferability of such a methodology.

Concerning the context of the SUNRISE project, the presented ex-ante methodology represents the first iteration of the SIA framework, aimed at analysing potential impacts on the development stage of the solutions. Therefore, as the presented approach is part of an ongoing SIA, where only the first round of assessment (ex-ante analysis) is completed, the overall results of the evaluation are not yet available. The findings from the initial phase of impact assessment lay the foundation for the next phases of the evaluation process: the subsequent assessment phase will occur following the pilot implementation of solutions, employing a mixed-method approach that focuses on the implications of toolset implementation for end-users. Due to time constraints, for the moment, only a limited number of focus groups were conducted, while additional focus groups, semi-structured interviews and surveys (engaging other stakeholder groups) are planned to be carried out at the subsequent stages of iterative SIA.

VI. CONCLUSION

This article outlines the rationale, strategy, and methodological phases of the ex-ante impact assessment, focusing on its application to the RII solution within the SUNRISE project. The evaluation strategy is designed to conduct an iterative social impact assessment of the technological solutions throughout the project's lifecycle. The ex-ante approach presented addresses the need for a standardized framework, facilitating a comprehensive understanding of potential social impacts and allowing for adjustments based on stakeholder input before deployment.

The article details the initial phase of a social impact assessment strategy, involving the definition and verification of impact categories specific to the RII solution. It begins with establishing a robust set of indicators and employs a participatory verification approach to identify critical subcategories of social impact. These phases aim to enhance understanding of the technological context, provide valuable feedback to developers, and optimize resource allocation for project-based social impact assessments prior to implementation.

ACKNOWLEDGMENT

This research has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101073821 (SUNRISE).

REFERENCES

- [1] F. Vanclay, "International Principles for Social Impact Assessment", *Impact Assessment and Project Appraisal*, vol. 21, no. 1, pp. 5-12, 2003, doi: 10.3152/147154603781766491
- [2] F. Vanclay, A.M. Esteves, I. Aucamp, and D. Franks, "Social Impact Assessment: Guidance for assessing and managing the social impacts of projects", 2015.
- [3] M. Di Domenico, H. Haugh, and P. Tracey, "Social bricolage: Theorizing social value creation in social enterprises", *Entrepreneurship theory and practice*, vol. 34, no. 4, pp. 681-703, 2010.
- [4] A. Fonseca, Ed. *Handbook of environmental impact assessment*. Edward Elgar Publishing, 2022
- [5] S. Kah and T. Akenroye, "Evaluation of social impact measurement tools and techniques: a systematic review of the literature", *Social Enterprise Journal*, vol. 16, no. 4, pp. 381-402, 2020.

- [6] D. Haski-Leventhal and A. Mehra, "Impact measurement in social enterprises: Australia and India", *Social Enterprise Journal*, vol. 12, no. 1, pp. 78-103, 2016.
- [7] O. Laedre, T. Haavaldsen, R.A. Bohne, J. Kallaos, and J. Lohne, "Determining sustainability impact assessment indicators", *Impact Assessment and Project Appraisal*, vol. 33, no. 2, pp. 98-107, 2015.
- [8] S.P. de Jong and R. Muhonen, "Who benefits from ex ante societal impact evaluation in the Euro-pean funding arena? A cross-country comparison of societal impact capacity in the social sciences and humanities", *Research Evaluation*, vol. 29, no. 1, pp. 22-33, 2020.
- [9] C. Pursiainen and E. Kytömaa, "From European critical infrastructure protection to the resilience of European critical entities: what does it mean?", *Sustainable and Resilient Infrastructure*, 8(sup1), pp. 85-101, 2023.
- [10] D. Haarsma and P.Y. Georgiadou, "Geo-ethics Requires Prudence with Private Data: GIM International interviews Professor Yola Georgiadou", *GIM Int.*, vol. 31, no. 10, pp. 16-19, 2017.
- [11] K.B. Culver, "From Battlefield to Newsroom: Ethical Implications of Drone Technology in Journalism", *Mass Media Ethics*, 29, pp. 52-64, 2014.
- [12] C.M. Gevaert, R. Sliuzas, C. Persello, and G. Vosselman, G., "Evaluating the societal impact of using drones to support urban upgrading projects", *ISPRS international journal of geo-information*, vol. 7, no. 3, p. 91, 2018.
- [13] R.L. Finn and D. Wright, D., "Privacy, data protection and ethics for civil drone practice: A survey of industry, regulators and civil society organisations", *Comput. Law Secur. Rev.*, 32, pp. 577-586, 2016.
- [14] C. Gagnon, P. Hirsch, and R. Howitt, "Can SIA empower communities?", *Environmental Impact Assessment Review*, 13(4), 1993, pp. 229-253.
- [15] R. Howitt, "Theoretical foundations". In *New Directions in Social Impact Assessment*, F. Vanclay and A.M. Esteves (Eds). Cheltenham; Northampton, MA: Edward Elgar Publishing, 2011, pp. 78-95.
- [16] A.J. Imperiale and F. Vanclay, "Using social impact assessment to strengthen community resilience in sustainable rural development in mountain areas", *Mountain Research & Development*, vol. 36, no. 4, 2016, pp. 431-442.
- [17] L.K. Mottee and R. Howitt, "Follow-up and social impact assessment (SIA) in urban transport-infrastructure projects: insights from the parramatta rail link", *Australian planner*, 55(1), pp. 46-56, 2018.
- [18] A.J. Imperiale and F. Vanclay, "From project-based to community-based social impact assessment: New social impact assessment pathways to build community resilience and enhance disaster risk reduction and climate action", *Current Sociology*, 2023, 00113921231203168.
- [19] R. Turner and R. Müller, "On the nature of project as a temporary organization", *International Journal of Project Management*, vol. 21, no. 1, pp. 1-8, 2003.
- [20] J.G. Turnley, "Social, cultural, economic impact assessments: A literature review. Prepared for The Office of Emergency and Remedial Response", *US Environmental Protection Agency*, 2002.
- [21] H. Bakker, R. Arkesteijn, M. Bosch-Rekvelde, and H. Mooi, "Project success from the perspective of owners and contractors in the process industry", in *The 24th IPMA World Congress*, Istanbul, Turkey, 1-3 November 2010, pp. 1-6.
- [22] E. Costa and C. Pesci, "Social impact measurement: why do stakeholders matter?", *Sustainability Accounting, Management and Policy Journal*, vol. 7, no. 1, pp. 99-124, 2016.
- [23] A. Passani, F. Monacciani, S. Van Der Graaf, F. Spagnoli, F. Bellini, M. Debicki, and P. Dini, "SEQUOIA: A methodology for the socio-economic impact assessment of Software-as-a-Service and Internet of Services research projects", *Research evaluation*, vol. 23, no. 2, pp. 133-149, 2014.
- [24] A. Pastra, T.M. Johansson, V. Alexandropoulou, N.L. Trivyza, and K. Kontaxaki, "Addressing the hazards of remote inspection techniques: a safety-net for vessel surveys", *Law, Innovation and Technology*, vol. 16, no. 1, pp. 43-76, 2024.
- [25] C. Sandbrook, "The social implications of using drones for biodiversity conservation", *Ambio*, 44(Suppl 4), pp. 636-647, nov 2015.
- [26] G.P. Jones, L.G. Pearlstone and H.F. Percival, "An assessment of small unmanned aerial vehicles for wildlife research", *Wildlife Society Bulletin*, 34, pp. 750-758, 2006.
- [27] H.-T. Lee, L.A. Meyn, and S. Kim, "Probabilistic safety assessment of unmanned aerial system operations", *Journal of Guidance, Control and Dynamics*, 36, pp. 610-617, 2013.
- [28] S. Kreps and J. Kaag, "The use of unmanned aerial vehicles in contemporary conflict: A legal and ethical analysis", *Polity*, 44, pp. 260-285, 2012.
- [29] J.D. Lewis and T. Nkuintchua, "Accessible technologies and FPIC: Independent monitoring with forest communities in Cameroon", *Participatory Learning and Action*, 65, pp. 151-165, 2012.
- [30] M. Yaqot and B. Menezes, "The good, the bad, and the ugly: review on the social impacts of unmanned aerial vehicles (UAVs)", in *International Conference of Reliable Information and Communication Technology*, Cham: Springer International Publishing, dec 2021, pp. 413-422.
- [31] L.L. Dhirani, N. Mukhtiar, B. S. Chowdhry, and T. Newe. "Ethical dilemmas and privacy issues in emerging technologies: A review", *Sensors*, vol. 23, no. 3, p. 1151, 2023.
- [32] V.A. Cloquell-Ballester, V.A. Cloquell-Ballester, R. Monterde-Diaz, and M.C. Santamarina-Siurana, "Indicators validation for the improvement of environmental and social impact quantitative assessment", *Environmental Impact Assessment Review*, vol. 26, no. 1, pp. 79-105, 2006.
- [33] S. Schuck-Zöllner, J. Cortekar, and D. Jacob, "Evaluating co-creation of knowledge: from quality criteria and indicators to methods", *Advances in Science and Research*, 14, pp. 305-312, 2017.
- [34] E. Brügger and P. Willems, "A critical comparison of offline focus groups, online focus groups and e-Delphi", *International Journal of Market Research*, vol. 51, no. 3, pp. 1-15, 2009.
- [35] M. Rainock, D. Everett, A. Pack, E.C. Dahlin, and C.A. Mattson, "The social impacts of products: a review", *Impact assessment and project appraisal*, vol. 36, no. 3, pp. 230-241, 2018.
- [36] M. Skare and M. Porada-Rochoń, "Technology and social equality in the United States", *Technological forecasting and social change*, 183, 121947, 2022.
- [37] T.C. Lindsey, "Sustainable principles: common values for achieving sustainability", *Journal of Cleaner Production*, vol. 19, no. 5, pp. 561-565, 2011.
- [38] A. Elmualim, R. Valle, and W. Kwawu, "Discerning policy and drivers for sustainable facilities management practice", *International Journal of Sustainable Built Environment*, 1(1), 16-25, 2012, doi: 10.1016/j.ijsbe.2012.03.001
- [39] U.J. Adama, *A study of the impact of technological innovations on the local sustainability of facilities management employees in South Africa*, 2019.
- [40] S. Amine and P.L. Dos Santos, P.L., "Technological choices and unemployment benefits in a matching model with heterogenous workers", *Journal of Economics*, vol. 101, no. 1, pp. 1-19, 2010.
- [41] G.L. Violante, "Skill-biased technical change", in *The New Palgrave Dictionary of Economics*, S.N. Darlauf and L.E. Blume (eds), Palgrave Macmillan, London, 2008.
- [42] G.L. Rafnsdottir and M.L. Gudmundsdottir, "New technology and its impact on well being", *Work*, vol. 22, no. 1, pp. 31-39, 2004.
- [43] M.J. Carley and E. Bustelo, *Social impact assessment and monitoring: a guide to the literature*. Routledge, 2019.
- [44] W.F. Cascio and R. Montealegre, "How technology is changing work and organizations", *Annual review of organizational psychology and organizational behavior*, vol. 3, pp. 349-375, 2016.
- [45] S.O. Hansson, "Theories and methods for the ethics of technology", *The ethics of technology: Methods and approaches*, pp. 1-14, mar 2017.
- [46] S. Chuang and C.M. Graham, "Embracing the sobering reality of technological influences on jobs, employment and human resource development: a systematic literature review", *European Journal of Training and Development*, vol. 42, no. 7/8, pp. 400-416, 2018.
- [47] J.D. Lee, "Review of a pivotal human factors article: 'humans and automation: use, misuse, disuse, abuse'", *Human Factors*, vol. 50, no. 3, pp. 404-410, 2008.
- [48] C. Panari, G. Lorenzi, and M.G. Mariani, "The predictive factors of new technology adoption, workers' well-being and absenteeism: the case of a public maritime company in Venice", *International Journal of Environmental Research and Public Health*, vol. 18, no. 23, 12358, 2021.
- [49] K.A. Pollard, B.T. Files, A.H. Oiknine, and B. Dalangin, "How to prepare for rapidly evolving technology: Focus on adaptability", in *Technical Report ARL-TR-9432*, US Combat Capabilities Development Command, Army Research Laboratory, Aberdeen Proving Ground United States, 2022.

LeAF: Leveraging Deep Learning for Agricultural Pest Detection and Classification for Farmers

Aditya Sengupta

Email: adityasngpta@gmail.com

Abstract—Farmers face many challenges while growing crops such as monitoring and maintaining plant health. Key indicators of poor plant health are plant anomalies such as pests, plant disease, and weeds, which can decrease crop yield. Over 40% of global crop production is lost to plant anomalies, costing \$220 billion annually. As the global population and demand for food increases, farmers will have to grow more food, making manual surveying for plant anomalies increasingly difficult. This forces farmers to excessively and indiscriminately apply more fertilizers and pesticides across their whole fields, often to both healthy and unhealthy plants, unnecessarily wasting acres worth of chemicals and increasing chemical contamination of food and environmental footprint of agriculture as the chemicals release greenhouse gases after their application and leak into ecosystems. Recent advances in deep learning with Convolutional Neural Networks (CNNs) allow using imaging data to solve this problem. LeAF aims to provide farmers with an end-to-end system to survey crops on the field and take targeted actions to maintain plant health. By focusing on agricultural pests, this paper demonstrates the following capabilities for the visual perception sub-system of LeAF: (1) use CNNs on field images to get plant-specific data with bounding box based detection and classification about plant anomalies at human-level accuracy and (2) combine detection and classification functionality into a single compact distilled model that can run on farmer accessible mobile phones or in embedded devices in agricultural tractors and robots with low latency and high throughput to enable real-time processing on video feeds. With lightweight and accurate plant anomaly detection and classification, LeAF addresses plant health management challenges faced by farmers, empowering them with actionable insights to enhance productivity while minimizing chemical usage and its environmental impact.

I. INTRODUCTION

Modern day agriculture has an enormous environmental footprint. In 2022, agriculture accounted for approximately 26% of global greenhouse gas emissions, used 50% of global inhabitable land, and accounted for 70% of all freshwater usage. [1]. This means the agriculture industry has one of the largest impacts on climate change. Over the next few decades, agriculture’s environmental footprint will only get larger because about 60% more food needs to be grown to feed the increasing world population [2]. This means immediate and effective steps need to be taken to increase the efficiency of agriculture to improve yields, conserve resources, and reduce the environmental impacts.

A. Problem

Farmers face challenges with continuously monitoring and maintaining plant health. Key indicators of poor plant health are plant anomalies such as pests, plant disease, and weeds,

which can decrease crop yield. Over 40% of global crop production is lost to plant anomalies, costing \$220 billion annually [3], [4].

Farmers often use chemicals to treat plant anomalies and ensure healthy crops, high yields, and good quality food for consumers. However, a lot of the emissions from agriculture come from the use of these chemicals (pesticides, herbicides, and fertilizers) to maintain plant health and treat plant anomalies (pests, weeds, and plant disease, respectively). Chemical-related greenhouse gas emissions such as nitrous oxide and methane can have up to a 300x higher global warming impact than carbon dioxide [5]. The chemicals emit greenhouse gases in their manufacturing process (creating harmful ground-level ozone), transportation to farms (fuel emissions from trucks), and even after they are applied to crops as they stimulate the production of nitrous oxide and methane in the soil. Furthermore, chemicals often leak into environments surrounding farms and into the ocean, harming wildlife and disrupting food chains.

These chemicals also have further downstream harmful impact. For example, on top of yield losses from pests, an additional \$60 billion dollars is spent on 1 billion pounds of pesticides annually in the US alone [6]. Furthermore, pesticide residues on food lead to 20,000+ new cases of cancer every year and contamination in nature from pesticides results in 80+ million fish and bird deaths annually [7]. However, with more infestations of invasive pests and the cultivation of more crops, the use of chemicals is only increasing [8].

The use of chemicals also increases production costs for farmers, increasing the cost of food for consumers. Not only are 78% of the world’s poor people farmers [9], their income is also declining because of falling crop prices since 2014 despite inflation [10]. Reducing chemical costs can help farmers keep up with decreasing crop sales revenue.

Most farmers are forced to apply chemicals across their whole field to both healthy and unhealthy crops because they have no data on where the pest infestations are. This wastes acres worth of chemicals. In order to reduce the use of chemicals, they need to be used more efficiently to still counteract plant anomalies and allow for healthy yield. A farmer could survey their crops themselves and analyze plant anomalies manually, but they might not have enough knowledge about the types of plant anomalies, whether they are helpful, benign, or malignant to the farm ecosystem, or which chemicals to use in order to treat them. In addition, many farmers have acres of land which makes self-surveying

unfeasible due to the shortage of time and labor to make accurate observations on a daily or weekly basis. With such a lack of data on plant anomalies in their field, farmers are less aware of the efficacy of their techniques, which may lead to bad yields despite the use of chemicals.

This highlights the need for automation in the agricultural industry to detect and identify plant anomalies and help farmers with treatment strategies that minimize the use of chemicals, cost, and harm to the environment. A solution which can provide farmers with such data will pave the way for better quality yields while using less chemicals in a more efficient and effective manner and minimizing costs and environmental impact.

B. Related Work

There is limited work on using machine learning to identify plant anomalies; these also have gaps that are addressed by LeAF. The IP102 pest classification dataset [11] is used in the same paper to train a classification model with accuracy significantly below human level (49.4% with ResNet-50). IP102 also contains more than 100 classes, which requires a larger model to achieve this accuracy while trading off inference latency. Since there are thousands of pests that need larger model to classify accurately, it is instead more effective to use a smaller, lightweight model tuned specifically for local pests, as in LeAF, since farmers only deal with a few major pests in their area.

For plant disease classification, AMAizeD focuses on corn diseases [12]. For weed classification, recent work is limited to weeds that impact bell peppers [13].

The problem with these models is that they have low accuracy and are not generalizable to different types of anomalies for other plants. For example, if you give AMAizeD a non-corn disease image, it will still classify it as a specific corn disease. This is especially dangerous for pests because the model might classify an unknown but helpful insect as a pest, leading to the farmer applying useless pesticides. Moreover, these models focus on classification, not detection, so they don't provide bounding boxes to pinpoint the anomalies. When dealing with multiple anomalies and surveying multiple crops, knowing where the anomaly is in the image is necessary for explainability, to keep a count of anomalies, and avoid both over and under counting.

LeAF, on the other hand, uses transfer learning to quickly customize for regional and local scenarios with small datasets, provides human level classification and detection accuracy for plant anomalies, classifies unseen pests as "unknown" to avoid incorrect treatment, and is also amenable to expanding supported types of anomalies with incremental retraining.

C. Goals & Benefits of LeAF

LeAF advances plant anomaly detection and mapping for farmers, offering an end-to-end solution that enhances efficiency and minimizes chemical usage. By leveraging emerging agricultural robots or tractors equipped with cameras, LeAF captures and analyzes crop images, initially focusing on pests

while designing to expand to diseases and weeds. By consolidating various plant anomaly detection tasks into a single lightweight model per anomaly type, the solution can run effectively on embedded devices in the edge, such as the EarthSense TerraSentia robot [14]. This streamlined approach not only ensures accuracy but also enables rapid tuning to different farming regions.

The ultimate objective of the LeAF end-to-end system [15] is to provide farmers with actionable insights, including treatment suggestions and cost-effectiveness estimates, through a user-friendly natural language based multimodal interface. By correlating plant anomaly data with mapped field arrays, LeAF enables farmers to monitor anomalies on a day-to-day basis and track the efficacy of their interventions, facilitating informed decision-making and optimizing crop yield.

To achieve this, LeAF utilizes precise yet lightweight model architectures like ResNet-18 [16] and YOLOv8 [17], trained on agricultural images sourced from the iNaturalist library [18]. The camera feed obtained while surveying the field (with robots or tractors, for example) is processed for anomaly detection and classification, with anomalies tracked with location awareness across the field. This individual plant level data is then mapped to a field array, facilitating analysis by an agriculture domain adapted Large Language Model (LLM) for treatment suggestions and Q&A support.

By empowering farmers with real-time insights and optimized chemical application strategies, LeAF aims to minimize environmental impact and maximizes crop yield, ultimately enhancing agricultural practices for a sustainable future. This paper covers the visual perception component of LeAF, involving plant anomaly detection and classification, and focuses on pests for concrete model development and evaluation.

II. PEST DETECTION AND CLASSIFICATION

The visual perception component of LeAF detects and classifies pests based on a given number of pest classes. The classes that LeAF aims to detect depend on the crop being grown and the geographical region of the world. For the prototype, I target ten main types of pests for corn and soybean crops grown in the midwestern region of the United States.

A. Initial Pipeline Overview

The solution comprises two stages: pest detection and classification. Initially, I develop a pipeline using foundation models for pest detection (with bounding box), evaluate its performance, and refine it. Subsequently, I focus on gathering datasets for training and evaluating a specialized pest classification model.

Given the challenge of obtaining large labeled datasets, I leverage foundation models like GroundingDINO [19] for pest detection. Although these models provide bounding boxes for pests, they lack fine-tuned classification accuracy for diverse agricultural pests. To address this, I utilize cropped images from GroundingDINO to feed into a specialized, more efficient classification model like ResNet-18. This two-stage pipeline, depicted in Figure 1, first localizes potential pests using

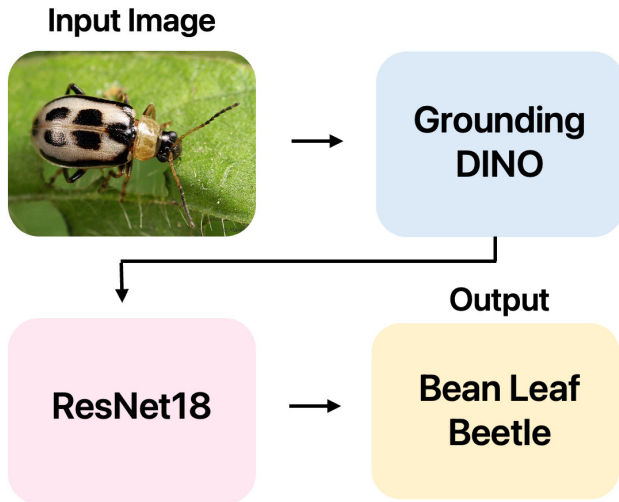


Fig. 1. Initial Pipeline with GroundingDINO and ResNet-18

GroundingDINO and then classifies them using ResNet-18. This approach mitigates data scarcity and enables recognition of unseen pest classes through an ‘unknown’ class.

Further exploration involves distilling this pipeline into a single, lightweight YOLOv8n model for mobile and edge deployment, as discussed in Section III.

B. Datasets & Training Prerequisites

After figuring out what the ML pipeline needs to input and output, building the initial pipeline, testing it, and coming up with a more refined solution, the next step is to train a custom model to classify pests. For training this and the eventual tweaking of the final pipeline, I first need to establish a dataset containing images of pests and their labels.

1) *Identifying Pest Classes*: LeAF focuses on the 10 most harmful pests affecting corn and soybean crops in the Midwestern United States, as identified by an UIUC entomology study [20]. These classes include Bean Leaf Beetle, Grape Colaspis, Japanese Beetle, Northern Corn Rootworm (CRW), Southern CRW, Western CRW, Grasshopper, Cloverworm/Looper, Stink Bug, and Dectes Stem Borer. This selection prioritizes the needs of farmers in this agriculturally significant region.

2) *Curating Initial Dataset*: Training machine learning models requires substantial data, which is often scarce for specific pest classes. LeAF leverages iNaturalist [18], a crowdsourced platform with hundreds of millions of images of various species, to curate a dataset of relevant pest images for model training.

The iNaturalist dataset includes images of all 10 of the above pest classes, so I created the training dataset by downloading images from iNaturalist. Utilizing iNaturalist’s export tool, I filtered images meeting research-grade standards and consensus thresholds. I ensured a balanced distribution, with around 1,000 images per class (total 10,000 images across

all classes), and split them into 80% for training, 10% for validation, and 10% for test.

3) *Device Platform*: For training and inference, I use Python 3.10 and the PyTorch ML framework on multiple hardware platforms, including Google Colab with Nvidia T4 GPUs, on-prem cluster of Nvidia A100 GPUs, and M1 Mac Mini desktop.

C. Training

With this pipeline structure of GroundingDINO feeding into a smaller CNN, it was now time to select the CNN model structure and custom train it on the pest dataset. Initially, I started with a ResNet-18 model [21], which is the smallest in ResNet family with one of the most simple CNN structures.

1) *Initial Training Process*: I initialized a ResNet-18 model pre-trained on ImageNet and replaced its last layer with a 10-class fully connected layer for pest classification [21]. Leveraging transfer learning, I utilized previously learned features to expedite training with less data. This approach also enables selective training of the final layer, conserving computational resources and accelerating training.

2) *Enhancing Accuracy*: Training ResNet-18 on the dataset for 100 epochs yielded a stagnant accuracy around 60%. Despite comparable training and validation accuracy, both were below human accuracy levels (which are above 90%), indicating underfitting. Experimenting with a larger model, ResNet-101, raised accuracy to nearly 80%, but extended training time and inference cost and latency beyond feasibility. Opting for computational efficiency, I retained the smaller ResNet-18 model and focused on the following optimizations to reach target accuracy of 90%+.

3) *Adjusting/Augmenting Data to Increase Accuracy*: After reassessing the pipeline based use case, I realized that ResNet-18 only needed to do inference on already cropped pest images from GroundingDINO’s output bounding boxes. Instead of the initial pest images, using cropped images addresses underfitting because the model will get less confused about the distracting image backgrounds and more focused on the pest and its features, since the image is essentially more zoomed in to the pest after cropping. Therefore, for the training and validation data, I fed all the images through GroundingDINO and cropped them to the bounding boxes for pests in the image. This made the input images for the ResNet-18 smaller and more focused, allowing the model to pick up on the features of the pest, therefore increasing accuracy.

Dataset	Train Accuracy	Validation Accuracy
Original	66.67%	65.71%
Cropped	82.60%	81.29%

TABLE I
RESNET-18 MODEL ACCURACY AFTER 10 EPOCHS OF TRAINING ON ORIGINAL AND CROPPED PEST DATASETS.

This change increased accuracy from 60% to 80% as shown in Table I. After this, I made smaller improvements to the training procedure by adjusting hyperparameters to further increase accuracy as follows.

4) *Adjusting Training Hyperparameters:* To further increase the accuracy from 80% to human levels (90%+) and speed up training, I adjusted many of the training hyperparameters for the ResNet-18 model. The adjustments in this section are incremental and cumulative. For example, if the optimizer was changed to Adam, then all future testing includes this change. In addition, all results were achieved by training for 10 epochs (unless otherwise specified). Training accuracy was evaluated based on correct/incorrect classifications on the training dataset and validation accuracy is based on the validation dataset. Here are the hyperparameters that were tuned:

Optimizer and Learning Rate: I started by using the Stochastic Gradient Descent (SGD) training optimizer. In addition to the other optimizations, this was giving about an 85% accuracy. However, when experimenting with other optimizers, I found the Adam Optimizer to work best. Next, I tuned the learning rate. While the learning rate is less important with the presence of optimizers which can adapt the learning rate based on the change in loss, it still plays an important role in providing an estimate or range. I found the perfect balance of decrease in loss to optimal end-accuracy at a learning rate of 0.001. Specifically, using the Adam optimizer helped the model jump to 91% accuracy, which now met the threshold of human performance. These optimizer results are shown in Table II and the learning rate results are shown in Table III.

Optimizer	Train Accuracy	Validation Accuracy
SGD	82.60%	81.29%
Adam	89.26%	88.97%

TABLE II

RESNET-18 MODEL ACCURACY USING SGD AND ADAM OPTIMIZERS.

Learning Rate	Train Accuracy	Validation Accuracy
0.01	89.26%	88.97%
0.001	91.60%	90.96%
0.0001	90.51%	90.18%

TABLE III

RESNET-18 MODEL ACCURACY USING DIFFERENT LEARNING RATES FOR ADAM OPTIMIZER.

Batch Size: Batch size is an important training hyperparameter to adjust. On one hand, increasing batch size improves training throughput (by increasing compute utilization) but requires more memory. On the other hand, increasing batch size may have negative effects on regularization, training convergence, and model quality. I found that a batch size of 32 was optimal as it maximized CPU and GPU utilization for the amount of RAM available, allowing for fast training. This batch size was also favorable for obtaining high accuracy. These results are shown in Table IV.

Layers to Train: As mentioned before, larger models like the ResNet-101 took more time and resources to train than smaller models like ResNet-18. One way to make the training process even faster is to train only the last (fully connected) layer of the model. This works because with a pre-trained




Fig. 2. Training and validation loss over 25 epochs. The validation loss flattens out after 10 epochs.

ResNet-18, the previous pre-training on ImageNet still helped tune the earlier layers to classify features present in any image like edges, contours, texture, and color. The only thing new is the classification of pests rather than everyday objects.

Training only the last layer decreases training time by 8x and produces almost identical accuracy results (within 1%), allowing to train for more epochs and tune other hyperparameters faster. The results are shown in Table V.

Number of Epochs: I found that 10 epochs were optimal for training this model by looking at the graphs plotting train and validation loss from 10 epochs (see Figure 2). Even though initial epochs allow bigger decreases in loss, this flattens out over time. After 10 epochs, the validation loss oscillates and does not decrease noticeably. While the training loss still decreases slowly, it is not necessary to train beyond this point because this means the model starts to overfit and make changes to weights that don't generalize well with new data.

5) *Summary:* Table VI summarizes the key milestones in the journey from 60% to 91%+ accuracy.

6) *Other Model Architectures:* I experimented with other model architectures like EfficientNet, MobileNet, ViT, and SqueezeNet. However, since this model is being deployed on the edge, it must be as lightweight as possible while still maintaining 90+% accuracy. Therefore, I had to decide on the tradeoff between accuracy and computational resources. The accuracy and inference times for the different models are in Table VII.

ResNet-18 exhibits the best efficiency-to-accuracy ratio among tested architectures, striking an optimal balance. Despite having more parameters than MobileNet, ResNet-18's performance speed remains comparable, indicating additional optimizations beyond parameter count. Post training, ResNet-18 achieved a 91.43% classification accuracy on the validation set comprising 1,000 cropped images.

In summary, the initial model pipeline (depicted in Figure 1) employs GroundingDINO to generate bounding box detec-

Batch Size	Average Training Time Per Epoch	Train Accuracy	Validation Accuracy
16	52.35 seconds	87.91%	87.56%
32	36.40 seconds	91.73%	91.43%
64	38.50 seconds	91.32%	90.87%

TABLE IV
RESNET-18 MODEL ACCURACY USING DIFFERENT BATCH SIZES.

Layer to Train	Train Accuracy	Validation Accuracy	Average Training Time Per Epoch
Train Only Last Layer	91.32%	90.87%	36.35 seconds
Train All Layers	91.45%	90.66%	273.46 seconds

TABLE V
RESNET-18 MODEL ACCURACY WHEN TRAINING LAST VS ALL LAYERS.

Tuning Technique	Validation Accuracy
Original Accuracy	65.71%
Data Filtering/Augmentation	81.29%
Adam Optimizer	88.97%
Learning Rate 0.001	90.96%
Batch Size 32	91.29%
Train 10 Epochs (Last Layer)	91.43%

TABLE VI

IMPACT OF TUNING TECHNIQUE ON VALIDATION ACCURACY (AFTER 10 EPOCHS OF TRAINING).

Model	Validation Accuracy	Inference Time per 1K images (sec)
ResNet-18	91.29%	36.35
MobileNetV2 S	87.65%	32.47
EfficientNetV2 L	93.32%	79.73
ViT B16	95.53%	126.82
SqueezeNet	89.96%	50.34

TABLE VII

PERFORMANCE OF DIFFERENT MODEL ARCHITECTURES.

tions, followed by ResNet-18 for pest classification, obtaining an accuracy of 90%+.

D. Example Detection Results

Figure 3 shows some sample results for detection and classification. The pipeline even performs well on tough examples like a brown grasshopper camouflaged in brown leaves, and a Corn Rootworm (CRW) camouflaged in a leaf and flower. In addition, for images where there is no pest, the pipeline works correctly because those images are filtered by the GroundingDINO model (which is why there wasn't a need to train a separate class for "no pests in image" scenarios on the ResNet-18).

III. DISTILLED LIGHTWEIGHT MODEL FOR MOBILE AND EDGE DEVICES

While the GroundingDINO and ResNet-18 based model pipeline has good accuracy, it is slow and takes 3 seconds per image on an Nvidia T4 GPU in Google Colab. To enable this to run in real time on an embedded device like a Raspberry Pi or Nvidia Jetson Nano, the pipeline needs to be shrunk drastically.

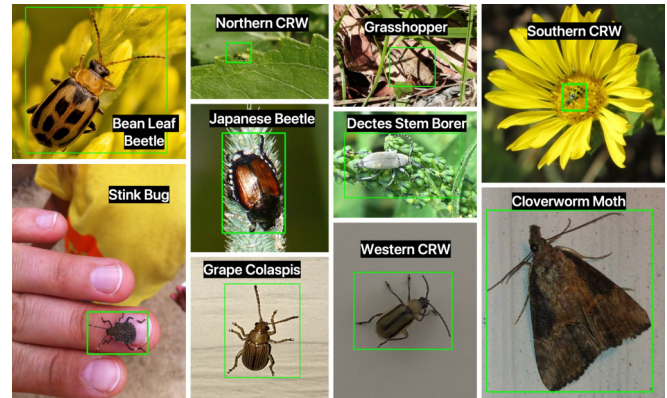


Fig. 3. Sample detection and classification results on test images with GroundingDINO and ResNet-18 pipeline

A. YOLOv8n instead of GroundingDINO and ResNet-18

The main bottleneck is the GroundingDINO model, which has 172M parameters compared to ResNet-18 which has only 11M parameters. One efficient model that can replace both is YOLO [17], which is an efficient one-shot computer vision model that can output bounding boxes and classifications. I will use YOLOv8n (Nano), latest lightweight version of YOLO with only 3.2M parameters (0.17% of the initial pipeline), which also enables huge increase in speed.

B. Model Distillation

Integrating YOLOv8n into the refined pipeline may challenge accuracy due to its smaller size. However, GroundingDINO's comprehensive bounding box training exceeds our need for pest classification alone. To transfer GroundingDINO's bounding box knowledge to YOLOv8n, I employ model distillation. This entails a larger foundation model annotating a task-specific dataset, from which a smaller student model learns. The previously annotated iNaturalist dataset using GroundingDINO, containing cropped images based on their bounding box coordinates to isolate pests with class labels, is used to train YOLOv8n. I used the PyTorch based Ultralytics API [22] to train the model for 10 epochs in about 30 minutes on a Mac Mini M1.




Fig. 4. Sample Detections and Classifications from YOLO Model (with confidence scores)

C. Evaluation on Pest Images

Testing the output of this YOLO model on some sample pest images yielded accurate bounding box results, shown in Figure 4. The model now takes only about 10 milliseconds per image on a mobile phone, compared to 3 seconds per image in initial pipeline on Nvidia T4 GPU, which is a 300x reduction in inference latency. This allows YOLO to perform real-time inference on video at 30 fps even on a mobile phone. The corresponding decrease in model size is 600x.

The model size and inference latency reduction benefits of the distilled model come with accuracy preservation comparable to that of the initial pipeline with GroundingDINO and ResNet-18 (as per the mAP metric). The distilled model has mAP50 of 0.81 (computed over $\text{IoU} \geq 50\%$). For the original pipeline, since GroundingDINO is a foundation model trained on vast amounts of data, it outputs perfect bounding boxes, so all of the detections satisfy the 50% overlap requirement of mAP50. Hence, the mAP calculation for initial pipeline simplifies to $mAP = \frac{1}{n} \sum_{i=1}^n \text{Precision}_i * \text{Recall}_i$ ($n = \text{\#classes}$). This gives an mAP of 0.82 for the original pipeline, hence mAP of the distilled model is very close, thus successfully preserving accuracy.

D. Handling Unknown Pest Classes

The initially trained YOLO struggled to classify out-of-distribution (OOD) pests, often misclassifying them as the most similar trained class. Detecting unknown pests becomes challenging, as class probabilities tend to dominate a single class rather than being spread across multiple classes.

Given the vast number of pest species globally, training YOLO to classify every pest is impractical, especially for edge deployment. I explored various techniques to detect unknown pests, ultimately settling on augmenting the original YOLO model with an ‘unknown’ pest class and trained it using a random distribution of pest images from the IP102 dataset.

This enhancement prevents misclassification of beneficial insects like ladybugs as harmful pests by assigning them to

the unknown class. Farmers can then manually review these images to determine their relevance. If numerous unknown pests emerge during deployment, the YOLO model can be extended to include additional pest classes, by incrementally expanding the last fully connected layer and retraining.

IV. CONCLUSION

LeAF offers lightweight, accurate pest detection and classification on mobile devices by distilling YOLO based model from Internet scale datasets using GroundingDINO foundational model for bounding box labeling and custom trained Resnet-18 model for classification. This visual perception capability for agricultural pests is part of the end-to-end LeAF system [15] that includes field level mapping for plant-by-plant analysis combined with LLM-powered analysis and recommendations to empower farmers with actionable insights for efficient plant health management. Informed by ongoing deployments and farmer feedback, LeAF is being further refined to ensure its effectiveness in promoting environmentally sustainable and productive agriculture.

REFERENCES

- [1] H. Ritchie, P. Rosado, and M. Roser. “Environmental impacts of food production,” *Our World in Data*, 2022.
- [2] J. G. D. Silva, “Feeding the world sustainably,” *United Nations Chronicle*, 2012.
- [3] “Invasive pest spread another fallout from climate change, UN-backed study finds,” *United Nations*, 2021.
- [4] N. Bhalla, “40% of global crop production is lost to pests. and it’s getting worse,” *World Economic Forum*, 2021.
- [5] J. Garthwaite, “Why laughing gas is a growing climate problem,” *Stanford News*, 2020.
- [6] S. LaMotte, “Reducing pesticides in food: Major food manufacturers earn an F grade,” *CNN*, 2023.
- [7] A. Pariona, “Top pesticide using countries,” *WorldAtlas*, 2017.
- [8] M. Tudi, H. D. Ruan, L. Wang, J. Lyu, R. Sadler, D. Connell, C. Chu, and D. T. Phung, “Agriculture development, pesticide application and its impact on the environment,” *National Library of Medicine*, 2021.
- [9] “For Up to 800 Million Rural Poor, a Strong World Bank Commitment to Agriculture,” *World Bank*, 2019.
- [10] W. A. Reinsch, T. Denamiel, and E. Kerstens, “Climate change and U.S. agricultural exports,” *CSIS*, 2023.
- [11] X. Wu and et al., “IP102: A Large-Scale Benchmark Dataset for Insect Pest Recognition,” *IEEE CVPR*, 2019.
- [12] A. Mall, S. Kabra, A. Lhila, and P. Ajmera, “AMaizeD: an end to end pipeline for automatic maize disease detection,” *ICST*, 2023.
- [13] A. Subeesh and et al., “Deep Convolutional Neural Network Models for Weed Detection in Polyhouse Grown Bell Peppers,” *Artificial Intelligence in Agriculture*, 2022.
- [14] “EarthSense TerraSentia,” <https://www.earthsense.co/robotics>.
- [15] A. Sengupta, “LeAF: Leveraging Deep Learning for Plant Anomaly Detection and Classification for Farmers with Large Language Models for Natural Language Interaction & BRANCH Robot-Based Deployment,” *IEEE CVPR Computer Vision for Science*, June 2024.
- [16] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” *IEEE CVPR*, 2016.
- [17] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” *IEEE CVPR*, 2016.
- [18] “iNaturalist,” <https://www.inaturalist.org/>.
- [19] S. Liu and et al., “Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection,” *arXiv.org*, 2024.
- [20] A. Decker and et al., “2022 applied research results field crop disease and insect management,” *UIUC Technical Report*, 2022.
- [21] “ResNet18 - Torchvision main documentation,” <https://pytorch.org/vision/main/models/generated/torchvision.models.resnet18.html>.
- [22] “Ultralytics YOLO,” <https://docs.ultralytics.com/>.

Key Factors Influencing Mobile Banking Adoption in Saudi Arabia

Amal Alzahrani
0009-0007-2568-1809
Dept. of Computer Science &
Information Systems, Baha University,
Baha, Saudi Arabia
Dept. of Informatics
University of Sussex,
Brighton, United Kingdom
Email: a.alzahrani@sussex.ac.uk

Natalia Beloff,
0000-0002-8872-7786
Dept. of Informatics
University of Sussex, Brighton,
United Kingdom
Email: n.beloff@sussex.ac.uk

Martin White
0000-0001-8686-2274
Dept. of Informatics
University of Sussex, Brighton,
United Kingdom
Email: m.white@sussex.ac.uk

Abstract—The introduction of mobile banking has revolutionized traditional financial practices, enhancing efficiency, customer experiences, and business models globally. Despite the global advancements in mobile banking, adoption rates remain low in Saudi Arabia. This paper seeks to identify key factors affecting adoption, using a mixed-methods approach. We propose a novel model integrating factors from the DeLone and McLean (D&M) model and the Unified Theory of Acceptance and Use of Technology (UTAUT2) model, complemented by additional factors. Data was gathered through online surveys and customer interviews. Findings revealed that net benefits, compatibility, facilitating conditions, and trust positively influence adoption, while literacy levels and digital skills pose barriers. Our study offers a significant theoretical contribution by synthesizing multiple models and enriches understanding of mobile banking adoption, aiding future research and industry decisions.

I. INTRODUCTION

THE banking sector has evolved significantly due to technological advancements, introducing various electronic channels. These innovations have transformed how banking services are accessed and utilized, reducing reliance on traditional methods and reshaping the industry's operational landscape. Among these innovations, mobile banking has emerged as a groundbreaking technology, offering customers unparalleled convenience and flexibility. Users prioritize mobility and 24/7 accessibility as the defining features of mobile banking, leading to increased adoption rates [1]. However, mobile banking usage remains lower than expected in Saudi Arabia, with only 59% of bank customers using mobile banking [2], compared to 84% in Kuwait [3], and 76% in Bahrain [4], highlighting the need to examine the factors influencing its adoption.

In Saudi Arabia, mobile systems have gained significant popularity, with the government emphasizing the importance of technology in enhancing the quality of life for Saudi citizens through Vision 2030. Despite these efforts, research indicates that Saudi mobile banking users are less satisfied compared to their UK counterparts [5]. Furthermore, half of Saudi banking customers have indicated they would consider switching to other banks due to insufficient mobile banking

features [6], raising concerns about the capability of these services to meet customer expectations.

Our study builds on [6] by focusing on specific inhibitors and facilitators of mobile banking adoption in Saudi Arabia, such as perceived trust, privacy, and compatibility. Using a mixed-method approach, we provide a detailed understanding of user behaviours with the latest data reflecting current trends. Our research also highlights unique inhibitors like literacy levels and digital skills, and facilitators such as the need for better educational resources, which were not covered previously.

This paper delves into the factors shaping mobile banking adoption in Saudi Arabia, crucial for enhancing customer satisfaction, fostering financial inclusion, and ensuring the success of mobile banking technologies. Guided by the research question, "What are the key factors influencing the adoption and usage of mobile banking in Saudi Arabia?", it identifies motivators and barriers by exploring technological, social, and individual factors. The study offers a tailored model for the Saudi context, providing insights to inform banks and policymakers on improving mobile banking services and encouraging broader adoption. This investigation aims to enrich the understanding of mobile banking and provide actionable recommendations for its advancement in Saudi Arabia.

II. LITERATURE REVIEW

This section reviews prior studies on mobile banking adoption to identify factors influencing its implementation in Saudi Arabia. It provides a comprehensive understanding of existing theories and empirical findings, informing the research model and hypotheses. Covering studies from the early 2000s to 2024, the literature review focuses on Saudi Arabia and includes relevant research from other Middle Eastern countries.

A. Research Framework

We utilize a theoretical framework that integrates factors from the DeLone and McLean information system success (D&M) model, proposed by [7], the Unified Theory of Acceptance and Use of Technology (UTAUT2), developed by [8], and additional factors, as depicted in Fig 1. This

integration aims to encompass a wide range of factors affecting mobile banking adoption and usage in Saudi Arabia, providing a holistic understanding of human, technological, and organizational influences on attitudes and behaviours towards mobile banking.

B. Methodology

We conducted a literature review using Google Scholar, PubMed, IEEE Xplore, ScienceDirect, and ProQuest, with keywords such as "mobile banking adoption," "UTAUT2," "D&M model," "customer satisfaction," "trust," and "privacy." Emphasizing relevance, recentness, and peer-reviewed publications, we screened abstracts and titles, followed by a full-text review, assessing quality based on research design, sample size, validity, reliability, and analysis robustness. This ensured the inclusion of high-quality studies, enhancing our understanding of mobile banking adoption in Saudi Arabia. Key factors influencing adoption are reported in our position paper [9], and repeated in section C.

C. Key Factors that influence adoption of mobile banking in Saudi Arabia

Performance Expectancy (PE) refers to the anticipated benefits from using a technology to improve performance [8]. Mobile banking offers convenient, flexible access to services anytime, anywhere, significantly impacting its adoption [6], [10]. Therefore, we expect customers to perceive mobile banking as enhancing overall performance.

Effort Expectancy (EE) refers to the perceived ease of using a new system [8]. Customers are more likely to adopt technology that is user-friendly [11]. Given the specific requirements of mobile banking and the need for certain levels of skills and knowledge, this factor can significantly influence customer intention to utilize it [10], supported by previous research [12]. Therefore, we expect customers to recognize the ease of use associated with mobile banking.

Social Influence (SI) refers to the impact of others' opinions on an individual's decision to adopt new technology [8]. The influence of social circles can significantly drive technology adoption [10], consistently highlighted in previous research [13]. Therefore, we expect customers to be motivated by others' preferences and views in using mobile banking.

Facilitating Conditions (FC) refer to the belief in adequate technical and organizational support for technology [8]. In mobile banking, customers are more likely to engage if there is sufficient support, skills, resources, and infrastructure [10], as highlighted by multiple studies [6], [12]. Therefore, we expect customers to recognize the presence of necessary infrastructure and resources for mobile banking usage.

Hedonic Motivation (HM) reflects the enjoyment and pleasure derived from using a technology [8]. It significantly impacts the adoption of new systems [13], with studies showing its strong influence on mobile banking use [6], [10].

Therefore, we expect customers to appreciate the pleasure and joy derived from using mobile banking.

System Quality (SQ) measures the effectiveness of technical services [7]. Research shows that a visually appealing and user-friendly mobile banking app enhances customer satisfaction [6], [14], [15]. Therefore, we expect customers to perceive mobile banking's ease of use, flexibility, functionality, and integration.

Service Quality (SVQ) measures the reliability, responsiveness, and timeliness of services [15]. Research shows that high service quality is crucial for mobile banking adoption [6], [16]. Therefore, we expect customers to perceive mobile banking reliability, availability, and competence.

Information Quality (IQ) refers to the relevance, accuracy, and adequacy of information provided by the system [15]. Mobile banking may face challenges due to limited screen sizes or app restrictions, impacting information access. Research shows that high information quality is crucial for user satisfaction [16]. Therefore, we expect customers to perceive the sufficiency, relevancy, and accuracy of information in mobile banking.

Perceived Privacy (PP) reflects confidence in data security during mobile banking usage, crucial for transaction privacy and risk mitigation [17]. Privacy concerns are recognized as a significant barrier to mobile banking adoption [18], reducing technology adoption and usage. Therefore, we expect customers to be assured that their information remains confidential and protected.

Perceived Trust (PT) reflects the belief in the alignment of actions with positive assumptions [19]. Establishing trust requires honesty, integrity, and benevolence within a system [20]. It significantly influences technology adoption by inversely relating to perceived risk [21]. Prior research robustly demonstrates the significant impact of this factor on mobile banking adoption [22], [23], [24]. Therefore, we expect customers to perceive the honesty and integrity of mobile banking.

Perceived Compatibility (PC) indicates how well an innovation aligns with user values, beliefs, and habits [25]. It holds significance as it notably impacts the acceptance of new technology [26]. It strongly influences technology adoption, as customers prefer innovations that fit their cultural beliefs [27]. Therefore, we expect mobile banking to be seen as compatible with customers' cultural values and beliefs.

Use of Mobile Banking involves evaluating customer actions and satisfaction with the system [7]. Understanding overall customer experiences is key to measuring satisfaction and usage rates [28]. Positive service experiences notably enhance user satisfaction [29], thereby increasing usage rates [30]. Therefore, we expect customers to perceive the advantages and benefits of mobile banking.

Satisfaction (SAT) measures customers' commitment and resistance to switching services, reflecting their loyalty [31]. Improving service quality is a key strategy for retaining customers and enhancing satisfaction [32]. Prior research

highlights the strong link between technology use and increased user satisfaction [6], [14]. Therefore, we expect customers to recognize the value and utility of mobile banking.

Loyalty (LOY) reflects a consistent use of the same service over time, showing commitment to it [33]. It is enhanced by improving service, information, and system quality [34]. Research underscores that user satisfaction significantly boosts loyalty to mobile banking [6]. Therefore, we expect customers to appreciate the value and utility of the service.

Net Benefits (BEN) evaluates the impact of technology on users or organizations, crucial for assessing effectiveness [35]. Positive net benefits increase the likelihood of adoption. Research shows that net benefits significantly affect mobile banking usage [27], and satisfaction [36]. Therefore, we expect customers to recognize and appreciate the advantages of mobile banking.

D. Research Model

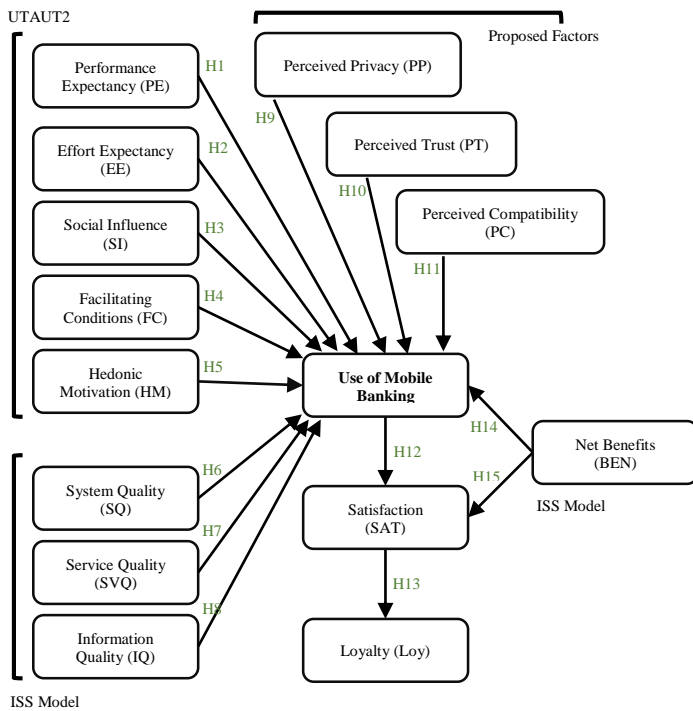


Fig 1. Research Model (adapted from UTAUT2 and D&M)

The theoretical framework combines established theories into a cohesive model to explore mobile banking adoption in Saudi Arabia. By analysing and empirically testing the relationships between key factors, the research enhances understanding of customer satisfaction and loyalty in mobile banking.

III. METHOD

This study employs a mixed-method approach to explore mobile banking adoption in Saudi Arabia. Data from 385 bank customers was collected via online surveys between January and April 2021. The survey, distributed through

email and WhatsApp, covered demographic details and mobile banking usage. Table I summarizes the participants' demographic characteristics.

TABLE I. DEMOGRAPHIC AND USAGE CHARACTERISTICS OF RESPONDENTS

		Frequency	Percent
Gender	Male	156	40.5
	Female	229	59.5
Age	18 - 24 years	134	34.8
	25 - 34 years	161	41.8
	35 - 44 years	68	17.7
	45 - 54 years	17	4.4
	55 or older	5	1.3
Marital Status	Married	150	39.0
	Single	224	58.2
	Divorced	11	2.9
Employment Status	Employed	133	34.5
	Self employed	54	14.0
	Unemployed	46	11.9
	Student	133	34.5
	Retired	3	0.8
Qualifications	Housewife/husband	16	4.2
	No schooling completed	1	0.3
	High school	37	9.6
	Diploma	7	1.8
	Bachelor	205	53.2
	Master or above	135	35.1

A. Quantitative Study

Analysis was conducted using Structural Equation Modeling SEM and Statistical Package for the Social Sciences (SPSS) software to infer relationships between variables. These tools are widely utilized in social sciences for such analyses.

• Convergent Validity and Model Fit

Convergent validity assesses the alignment among measurements of a specific construct [37], providing evidence of reliability and accuracy. One method to assess convergent validity is through calculating the Average Variance Extracted (AVE) for each construct. An acceptable AVE value, typically 0.5 or higher, signifies that the construct explains at least 50% of the variance in its items. Model fit refers to how well a statistical model fits the observed data. It is evaluated using multiple indices: the chi-square value should be less than 3 (<3), the standardized root mean square residual (SRMR) should be equal to or less than 0.10 (SRMR ≤ 0.10), the CMIN/DF value should be less than or equal to 5 (≤ 5), the comparative fit index (CFI) should be greater than or equal to 0.90 (CFI ≥ 0.90), and the root mean square error of approximation (RMSEA) should be equal to or less than 0.08 (RMSEA ≤ 0.08) [38], [39]. Tables II and III present the values of CMIN/DF, CFI, IFI, SRMR, RMSEA, and AVE, for evaluating model fit and convergent validity.

TABLE II. MODEL FIT INDICES

	Threshold	Value	Interpretation
CMIN/DF	≤5	2.1	Good fit
CFI	≥ 0.90	0.902	Good fit
IFI	≥ 0.90	0.903	Good fit
SRMR	≤ 0.10	0.043	Good fit
RMSEA	≤ 0.08	0.056	Good fit

TABLE III.
AVERAGE VARIANCE EXTRACTED VALUES

	No of Items	AVE
PE	2	0.60
EE	3	0.59
SI	2	0.75
HM	2	0.80
FC	2	0.54
SQ	6	0.51
IQ	5	0.53
PT	4	0.67
SQ	3	0.81
PP	4	0.57
PC	4	0.56
USE	3	0.62
SAT	3	0.68
LOY	3	0.55
BEN	3	0.55

The analysis revealed a chi-square value of 2237.823 with 1019 degrees of freedom and a probability value of < 0.001 . The CMIN/DF value was 2.1, indicating satisfactory fit. Fit measures, including CFI = 0.902, IFI = 0.903, SRMR = 0.043, and RMSEA = 0.056, all demonstrated satisfactory levels, with AVE values meeting the required standard for good model fit. Despite the chi-square value not indicating model fit, the study achieved five to six indices meeting the criteria, ensuring the model's overall fit.

- Reliability

Reliability measures the consistency of variable measurements [37]. It is assessed using Composite Reliability (CR) and Cronbach's alpha. In this study, Cronbach's alpha needed to be greater than 0.60. Values between 0.60 and 0.70 are satisfactory, while values of 0.80 or higher are very good [40]. The evaluation outcomes for reliability are presented in Table IV below.

TABLE IV.
RELIABILITY AND CONSTRUCT VALIDITY

	Cronbach's α	CR	Reliability and Construct Validity
PT	0.88	0.89	Yes
PP	0.85	0.84	Yes
PC	0.82	0.83	Yes
PE	0.74	0.75	Yes
EE	0.81	0.81	Yes
SI	0.85	0.85	Yes
FC	0.69	0.70	Yes
HM	0.88	0.89	Yes
SQ	0.86	0.86	Yes
IQ	0.85	0.85	Yes
SVQ	0.92	0.92	Yes
SAT	0.86	0.86	Yes
LOY	0.72	0.77	Yes
BEN	0.77	0.78	Yes
USE	0.83	0.83	Yes

As shown in the table, all items in the proposed model meet the minimum reliability requirement and exhibit satisfactory measurement properties.

- Discriminant Validity

We utilized the HTMT method, which is believed to be more appropriate method, to evaluate discriminant validity. It was argued that the HTMT value ought to be less than 0.90, as values exceeding this threshold indicate a lack of discriminant validity [41]. The findings of the discriminant validity analysis are shown in Table V below.

TABLE V.
DISCRIMINANT VALIDITY

	USE	BEN	LOY	SAT	SVQ	IQ	SQ	HM	FC	SI	EE	PE	PC	PP	PT
USE	-														
BEN	0.81	-													
LOY	0.77	0.84	-												
SAT	0.66	0.75	0.84	-											
SVQ	0.25	0.33	0.33	0.42	-										
IQ	0.52	0.63	0.56	0.65	0.56	-									
SQ	0.49	0.65	0.60	0.70	0.52	0.73	-								
HM	0.20	0.33	0.33	0.45	0.35	0.48	0.61	-							
FC	0.69	0.66	0.54	0.47	0.28	0.60	0.68	0.38	-						
SI	0.20	0.32	0.26	0.29	0.25	0.44	0.46	0.39	0.35	-					
EE	0.39	0.55	0.50	0.54	0.37	0.48	0.59	0.41	0.64	0.31	-				
PE	0.28	0.42	0.37	0.41	0.21	0.36	0.42	0.41	0.35	0.43	0.39	-			
PC	0.63	0.74	0.77	0.83	0.35	0.77	0.72	0.49	0.56	0.32	0.49	0.35	-		
PP	0.51	0.57	0.47	0.51	0.48	0.68	0.55	0.32	0.49	0.30	0.39	0.16	0.64	-	
PT	0.56	0.62	0.65	0.69	0.52	0.70	0.57	0.43	0.50	0.23	0.45	0.44	0.73	0.82	-

The table reveals that all constructs have HTMT values below 0.90, indicating that this study achieved sufficient levels of discriminant validity.

- Structural Model and Hypotheses Testing

The path coefficient evaluates hypotheses and indicates relationships between variables. A value close to 1 shows a strong positive correlation, while a value near -1 indicates a strong negative correlation [42]. Coefficients near zero are usually insignificant. Table VI summarizes the path coefficient analysis, showing the relationships between dependent and independent variables.

TABLE VI.
RESEARCH MODEL EVALUATION

	Path	Beta	t	p
H1	PE \rightarrow Use	0.212	4.256	0.000
H2	EE \rightarrow Use	0.322	6.651	0.000
H3	SI \rightarrow Use	0.172	3.420	0.001
H4	FC \rightarrow Use	0.518	11.837	0.000
H5	HM \rightarrow Use	0.173	3.434	0.001
H6	SQ \rightarrow Use	0.426	9.219	0.000
H7	SVQ \rightarrow Use	0.229	4.600	0.000
H8	IQ \rightarrow Use	0.449	9.834	0.000
H9	PP \rightarrow Use	0.463	10.210	0.000
H10	PT \rightarrow Use	0.488	10.941	0.000
H11	PC \rightarrow Use	0.532	12.305	0.000
H12	Use \rightarrow Sat	0.569	13.524	0.000
H13	Sat \rightarrow Loy	0.707	19.566	0.000
H14	Ben \rightarrow Use	0.644	19.543	0.000
H15	Ben \rightarrow Sat	0.618	15.384	0.000

As shown above, there are significant relationships among all path coefficients. Ben and Use had the highest significant positive path relationship. The relationship between PC and Use was shown to be the second most significant positive relationship, followed by the relationship between PT and Use. Among significant positive relationships, the least but significant positive relationship was between SI and Use. In summary, all paths showed significant relationships explaining approximately 62% of the variance in the use of mobile banking.

IV. DISCUSSION

The study investigated factors influencing mobile banking adoption in Saudi Arabia and confirmed all hypotheses as significant. Performance Expectancy was shown to significantly impact usage, indicating that customers are more likely to adopt mobile banking when they perceive its benefits in enhancing their performance. Effort Expectancy was also significant, with customers preferring technology that is convenient and user-friendly. Social Influence had the least positive impact, possibly due to Saudi Arabia's cultural norms emphasizing privacy and conservatism, which reduce the influence of social networks on adoption. Facilitating Conditions showed a significant positive impact, as users are more likely to adopt technology when it provides support, training, and guidance. Hedonic Motivation was confirmed to impact usage, though practical benefits like convenience and efficiency are the primary drivers in Saudi Arabia.

System Quality was significant, highlighting the need for reliable, functional, and user-friendly apps compatible with various devices, especially for older users. Service Quality also had a significant impact, boosting customer trust and satisfaction, with improved educational resources and training enhancing the user experience. Information Quality was crucial, with effective usage influenced by the accuracy, timeliness, and format of information, and impacted by literacy and language barriers. Perceived Privacy and Perceived Trust significantly influenced adoption, as users are more likely to engage with mobile banking when they trust that their data is secure and the service reliable. Perceived Compatibility eased adoption by aligning with existing practices and reducing resistance.

The relationship between mobile banking usage and user satisfaction was confirmed, with users valuing the efficiency, immediacy, and accessibility of mobile banking. User satisfaction and loyalty were also positively correlated, with positive experiences fostering trust and loyalty. Net Benefits had the strongest impact on adoption, suggesting individuals are inclined to adopt technology when the perceived benefits outweigh the costs and risks.

This paper extends previous research [6], by exploring deeper into factors such as perceived trust, privacy, and compatibility, using a mixed-method approach and recent data to offer a nuanced understanding of mobile banking adoption, incorporating new inhibitors and facilitators not previously addressed.

V. CONCLUSION AND IMPLICATIONS

This paper developed a framework to study mobile banking adoption in Saudi Arabia using a mixed-methods approach. Findings reveal that net benefits and perceived compatibility drive adoption, while literacy levels and digital skills act as barriers. The study fills a gap in understanding mobile banking adoption, offering insights that help banks refine strategies and contribute to academic and practical knowledge. It provides new quantitative data for future research and policy, and offers practical guidance for banks, service providers, and app developers to enhance user experience and satisfaction.

VI. RECOMMENDATION AND FUTURE WORK

The survey, based on convenience sampling, mostly involved young and educated respondents, limiting its generalizability. Efforts to include a more diverse age range were constrained, particularly by online data collection challenges during COVID-19. Future research should aim for a more representative sample across all demographics and regions to improve generalizability and insights for mobile banking providers.

Future research should compare regions to explore cultural impacts on mobile banking. Collaborative international studies could provide deeper insights into cross-cultural dynamics.

Given the evolving nature of mobile banking, a longitudinal approach could track changes in customer attitudes over time. Future research should explore technological, cultural, and contextual factors affecting mobile banking in Saudi Arabia and consider alternative models or theories to enhance understanding. Expanding the study to other technological services could further validate the model and offer broader insights.

REFERENCES

- [1] S. Raza, A. Umer, and N. Shah, "New determinants of ease of use and perceived usefulness for mobile banking adoption," *International Journal of Electronic Customer Relationship Management*, 11, pp. 44-65, 2017, <https://doi.org/10.1504/IJECRM.2017.086751>.
- [2] M. Jobeily, and M. Minawi, "Views On Digital Banking in KSA", *Ipsos*, 2021. Available at: <https://www.ipsos.com/sites/default/files/ct/news/documents/2021-10/Views%20on%20Digital%20Banking%20in%20KSA.pdf>.
- [3] N. Qahoush, and O. Fahmy, "Digital Adoption during the Pandemic", *Ipsos*, 2021. Available at: <https://www.ipsos.com/sites/default/files/ct/news/documents/2021-05/Spotlight.Kuwait%20-%20Digital%20Adoption%20During%20Pandemic.pdf>.
- [4] A.A. Alawadhi, "Adoption and acceptance of m-banking system in Kingdom of Bahrain", *International Journal on Information and Communication Technology (IJoICT)*, vol. 3, no. 2, 2021.
- [5] S. Al-Otaibi, N.R. Aljohani, M.R. Hoque, and F.S. Alotaibi, "The satisfaction of Saudi customers toward mobile banking in Saudi Arabia and the United Kingdom," *Journal of Global Information Management*, vol. 26, no. 1, pp. 85-103, 2018, DOI: 10.4018/JGIM.2018010105.
- [6] A.M. Baabdullah, A.A. Alalwan, N.P. Rana, H. Kizgin, and P. Patil, "Consumer use of mobile banking (M-banking) in Saudi Arabia: Towards an integrated model," *International Journal of Information Management*, vol. 44, pp. 38-52, 2019, <https://doi.org/10.1016/j.ijinfomgt.2018.09.002>.
- [7] W.H. DeLone and E.R. McLean, "The DeLone and McLean model of information systems success: A ten-year update," *Journal of*

- Management and Information Systems*, vol. 19, no. 4, pp. 9–30, 2003, DOI:10.1080/07421222.2003.11045748.
- [8] V. Venkatesh, M.G. Morris, G.B. Davis, and F.D. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quarterly*, vol. 27, no. 3, pp. 425–478, 2003, <https://doi.org/10.2307/30036540>.
- [9] Alzahrani, A., Beloff, N. and White, M. 'IMMBA – An Integrative Model for Mobile Banking Adoption: The Case of Saudi Arabia', *2021 3rd International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, 2021, DOI: 10.1109/HORA52670.2021.9461290.
- [10] A.A. Alalwan, Y.K. Dwivedi, N.P. P. Rana, and M.D. Williams, "Consumer adoption of mobile banking in Jordan: Examining the role of usefulness, ease of use, perceived risk and self-efficacy," *Journal of Enterprise Information Management*, vol. 29, no. 1, pp. 118–139, 2016, <https://doi.org/10.1108/JEIM-04-2015-0035>.
- [11] Y.K. Dwivedi, N.P. Rana, A. Jeyaraj, M. Clement, and M.D. Williams, "Re-examining the Unified Theory of Acceptance and Use of Technology (UTAUT): Towards a Revised Theoretical Model," *Information Systems Frontiers*, vol. 21, no. 3, pp. 719–734, 2019, DOI:10.1007/s10796-017-9774-y.
- [12] N.S. Alshamrani, *Trust as an influencer of the intention to use mobile banking smartphone application in Saudi Arabia*. Doctoral thesis (PhD), Metropolitan University, United Kingdom, 2018.
- [13] Y. Chetoui, H. Lebdaoui, and N. Hafid, "Mobile banking usage in the post pandemic era: Demystifying the disparities among divergent user segments in a majority-Muslim country", *Journal of Islamic Marketing*, 2023, DOI: 10.1108/JIMA-01-2022-0032.
- [14] O. Bouhlel, K. Garrouch, and M. Nabil, "Assessing the Success of Mobile Banking in Saudi Arabia: Re-Specification and Validation of the DeLone and McLean Model," *International Journal of E-Services and Mobile Applications*, vol. 15, no. 1, pp. 1-24, 2023, DOI: 10.4018/IJESMA.318088.
- [15] T. Zhou, "An empirical examination of continuance intention of mobile payment services," *Decision Support Systems*, vol. 54, no. 2, pp. 1085–1091, 2013, DOI:10.1016/j.dss.2012.10.034.
- [16] M.O. Mansour, "Acceptance of mobile banking in Islamic banks: integration of DeLone and McLean IS model and unified theory of acceptance and use of technology," *International Journal of Business Excellence*, vol. 21, no. 4, 2020, DOI:10.1504/IJBEX.2019.10021858.
- [17] Y.S. Wang, Y.M. Wang, H.H. Lin, and T.I. Tang, "Determinants of user acceptance of Internet banking: An empirical study," *International Journal of Service Industry Management*, vol. 14, no. 5, pp. 501–519, 2003, <https://doi.org/10.1108/09564230310500192>.
- [18] M. Merhi, K. Hone, A. Tarhini, "Cross-cultural study of the intention to use mobile banking between Lebanese and British consumers: Extending UTAUT2 with security, privacy and trust," *Technology in Society*, vol. 59, pp. 101-151, 2019, DOI:10.1016/j.techsoc.2019.101151.
- [19] Y. Köksal and S. Penez, "An investigation of the important factors influence web trust in online shopping," *Journal of Marketing Management*, vol. 6, no. 1, pp. 28-40, 2015.
- [20] D. Gefen, E. Karahanna, and D. W. Straub, "Trust and tam in online shopping: AN integrated model," *MIS Quarterly*, vol. 27, no. 1, pp. 51–90, 2003, DOI:10.2307/30036519.
- [21] S.K. Sharma and M. Sharma, "Examining the role of trust and quality dimensions in the actual usage of mobile banking services: An empirical investigation," *International Journal of Information Management*, vol. 44, October 2018, pp. 65–75, 2019, DOI: 10.1016/j.ijinfomgt.2018.09.013.
- [22] M. Al-Husein, and M.A. Sadi, "Preference on the perception of mobile banking: A Saudi Arabian perspective", *European Online Journal of Natural and Social Sciences*, vol. 4, no. 1, pp. 161, 2015.
- [23] M. Mansour, "Acceptance of Mobile Banking in Islamic Banks: Integration of DeLone & McLean IS model and Unified Theory of Acceptance and Use of Technology," *International Journal of Business Excellence*, 2019, DOI:10.1504/IJBEX.2019.10021858.
- [24] M. El-Masri and A. Tarhini, "Erratum to: Factors affecting the adoption of e-learning systems in Qatar and USA: Extending the Unified Theory of Acceptance and Use of Technology 2 (UTAUT2)," *Educational Technology Research and Development*, vol. 65, no. 3, pp. 765–767, 2017, DOI:10.1007/s11423-017-9526-1.
- [25] M.S. Sohail and I.M. Al-Jabri, "Attitudes towards mobile banking: Are there any differences between users and non-users?" *Behaviour and Information Technology*, vol. 33, no. 4, pp. 335–344, 2014, DOI:10.1080/0144929X.2013.763861.
- [26] D. Chen, B. Vallespir, and N. Daclin, "An approach for enterprise interoperability measurement," *CEUR Workshop Proceedings*, vol. 341, pp. 1–12, 2008.
- [27] M.S. Alzaidi, "Exploring the Determinants of Mobile Banking Adoption in the Context of Saudi Arabia," *International Journal of Customer Relationship Marketing and Management (IJCRM)*, vol. 3, no. 1, 2022, DOI: 10.4018/IJCRM.289206.
- [28] Y. K. Dwivedi, K. K. Kapoor, M. D. Williams, and J. Williams, "RFID systems in libraries: An empirical examination of factors affecting system use and user satisfaction," *International Journal of Information Management*, vol. 33, no. 2, pp. 367–377, 2013, <https://doi.org/10.1016/j.ijinfomgt.2012.10.008>.
- [29] S. Laforet and X. Li, "Consumers' attitudes towards online and mobile banking in China," *International Journal of Bank Marketing*, vol. 23, no. 5, pp. 362–380, 2005, DOI:10.1108/02652320510629250.
- [30] C. Tam and T. Oliveira, "Understanding mobile banking individual performance: The DeLone & McLean model and the moderating effects of individual culture," *Internet Research*, vol. 27, no. 3, pp. 538–562, 2017, DOI:10.1108/IntR-05-2016-0117.
- [31] P. Oppong and H. Adjei, "The Role of Information Technology in Building Customer Loyalty in Banking: (a Case Study of Agricultural Development Bank Ltd., Sunyani)," *British Journal of Marketing Studies*, vol. 2, no. 4, pp. 9–29, 2014, DOI:10.13140/RG.2.2.31181.46561.
- [32] J. Lee, J. Lee, and L. Feick, "The impact of switching costs on the customer satisfaction-loyalty link: Mobile phone service in France," *Journal of Services Marketing*, vol. 15, no. 1, pp. 35–48, 2001, <https://doi.org/10.1108/08876040110381463>.
- [33] Wang, S. *Factors Impacting the Uptake of Mobile Banking in China: Integrating UTAUT, TTF and ECM Models*. PhD thesis, Manchester University, United Kingdom, 2018.
- [34] Z. Saleem, K. and Rashid, "Relationship between customer satisfaction and mobile banking adoption in Pakistan," *International Journal of Trade, Economics and Finance*, vol. 2, no.6, 2011, pp- 537-544, DOI:10.7763/IJTEF.2011.V2.162.
- [35] N. Urbach and B. Müller, "The updated Delone and Mclean model of information systems success", in Dwivedi, Y.K., Wade, M.R. and Schneberger, S.L. (Eds.), "Information systems theory: Explaining and predicting our digital society", vol. 28, pp. 1–18, 2012, DOI: 10.1007/978-1-4419-6108-2_1.
- [36] H.B. Abdennebi, "M-banking adoption from the developing countries perspective: A mediated model," *Digital Business*, vol. 3, no. 2, 2023, <https://doi.org/10.1016/j.digbus.2023.100065>.
- [37] J.F. Hair, W.C. Black, B.J. Babin, and R.E. Anderson, *Multivariate data analysis: A global perspective*. New Jersey: Pearson Education Inc, 2010.
- [38] L. Hu, and P.M. Bentler, "Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives," *Structural Equation Modelling: A Multidisciplinary Journal*, 6, pp. 1–55, 1999, <http://dx.doi.org/10.1080/10705519909540118>.
- [39] L. Klem, "Structural equation modelling. In Grimm L.G. & Yarnold P.R. (Eds.), Reading and understanding MORE multivariate statistics," *American Psychological Association*, pp. 227–260, 2000.
- [40] C. Hulin, R. Netemeyer, and R. Cudeck, "Can a Reliability Coefficient Be Too High?," *Journal of Consumer Psychology*, 10, pp. 55-58, 2001, DOI:10.2307/1480474.
- [41] J. Henseler, C.M. Ringle, and M. Sarstedt, "A new criterion for assessing discriminant validity in variance-based structural equation modelling," *Journal of the Academy of Marketing Science*, 43, pp. 115–135, 2015, DOI:10.1007/s11747-014-0403-8.
- [42] M.S. Shahbaz, A.F. Chandio, M. Oad, A. Ahmed, R. Ullah, "Stakeholders' management approaches in construction supply chain: a new perspective of stakeholder's theory", *International Journal of Sustainable Construction Engineering and Technology*, Vol. 9, No. 2, pp. 16–25, 2018, DOI:10.30880/ijscet.2018.09.02.002.

A novel ensemble learning technique of shallow models applied on a COVID-19 dataset

D Babuc

ORCID: 0009-0000-5126-6480

Computer Science Department, West University of Timișoara

Blvd. Vasile Pârvan 4, 300223 Timișoara, Romania

Email: diogen.babuc00@e-uvt.ro

Abstract—Our lives were affected by the COVID-19 pandemic. In order to face this crisis, we provided a novel ensemble learning strategy to tackle the COVID-19 prediction and classification problems. Because of their capacity to handle the complex and varied nature of COVID-19 data, a range of shallow models, including K-Nearest Neighbors, Decision Trees, Support Vector Machines, Classification and Regression Trees, and Extreme Gradient Boost, are included in our method. Using a COVID-19 dataset, each model is trained independently and then ensemble learning techniques are used to integrate the predictions of the models. We use strict model validation and hyperparameter optimization to improve performance. Comparing our ensemble method to a single model or traditional ensemble techniques, our results show considerable improvements in classification performance and prediction accuracy.

Index Terms—Ensemble Learning, Machine Learning, COVID-19, Performance Metrics, Prediction and Classification.

I. INTRODUCTION

SINCE its appearance in late 2019, the COVID-19 pandemic has had an influence on cultures, economy, and healthcare systems across the globe [1]. Predictive modeling has become an essential process for studying and projecting the trajectory of the virus as governments and health organizations struggle to stop its spread [2]. In this paper, we investigate the creation of models that forecast the total number of COVID-19 cases worldwide. We used a dataset that runs through September 2020. Additionally, we classify nations into those with and without a higher risk of contracting SARS-CoV-2. Due to the COVID-19 pandemic’s intricacy, new methods of data analysis and forecasting have been required.

Our primary focus lies in exploring the predictive potential of historical data up to September 2020. We want to capture critical phases of the pandemic’s evolution. Through retrospective analysis, we aim to elucidate patterns, trends, and underlying factors influencing the spread of COVID-19 across different regions and timeframes. Using statistical indicators and machine learning techniques, we seek to construct predictive models capable of discerning the complex interplay between various epidemiological variables and forecasting the total cases of COVID-19 with precision and reliability. Through this interdisciplinary effort, we aim to contribute to ongoing

This work was supported by West University of Timișoara.

global efforts to combat the COVID-19 pandemic. Our goal is to provide stakeholders with the knowledge and resources they need to effectively navigate the obstacles presented by this unprecedented public health catastrophe by using the power of prospective and retrospective predictive modeling and data-driven insights.

We hope to provide a better understanding of the dynamics of the pandemic and enable informed decision-making in the face of uncertainty (Fig. 1).




Fig. 1. Objectives for COVID-19 infection risk estimation.

Concretely, our main objectives are:

- 1) Developing and optimizing predictive shallow machine learning models using historical worldwide COVID-19 data up to September 2020.
- 2) Building an ensemble learning model, called *Reเนสansa*, for investigating the impact of epidemiological factors, including active cases, total tests conducted, and population demographics, on the total number of COVID-19 cases worldwide.
- 3) Evaluating performance of the predictive models developed through evaluation metrics and statistical indicators.

We compare the models’ results with the observed data to gauge the trustworthiness and effectiveness of the selected forecasting and classification methodologies.

II. BACKGROUND INFORMATION AND RELATED WORKS

In this section, we will discuss the background information and previous studies that have explored various methodologies, from traditional statistical models to modern ensemble learning approaches. We want to accurately forecast transmission trends and classify disease outcomes.

A. COVID-19 Disease

Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) is a new betacoronavirus that is a member of the Coronaviridae family and the source of COVID-19 [3]. SARS-CoV-2 is a single-stranded, enveloped, positive-sense RNA virus with spike (S), envelope, membrane and nucleocapsid proteins. Viral entrance into host cells is mediated by the S protein, which binds to the angiotensin converting enzyme 2 receptor and promotes membrane fusion and viral multiplication.

A wide range of clinical signs are displayed by COVID-19, from moderate or asymptomatic sickness to severe respiratory failure [4]. Common symptoms include fever, cough, fatigue, and gastrointestinal symptoms. Severe cases are characterized by acute respiratory distress syndrome, multiorgan dysfunction, and thrombotic complications. Certain population groups, such as older adults and immunocompromised individuals, are at increased risk of severe disease and adverse outcomes.

B. COVID-19 Pandemics' Social Impact

One of the worst global health emergencies in recent memory, the COVID-19 epidemic has had a tremendous effect on civilizations all around the world. [5]. Since the new coronavirus SARS-CoV-2 first appeared in late 2019, the pandemic has quickly expanded throughout continents, overcoming geographic barriers and igniting hitherto unheard-of public health measures [6].

COVID-19 has challenged our understanding of infectious diseases and highlighted the interconnectedness of our modern world. From the outset, the pandemic has posed multifaceted challenges, ranging from containment efforts and healthcare delivery to social distancing measures and economic stability [7]. The COVID-19 pandemic has underscored the importance of rapid and coordinated responses from governments, healthcare institutions, and communities to mitigate transmission, protect vulnerable populations, and minimize the burden on healthcare infrastructure. Measures such as lockdowns, travel restrictions, mass testing, contact tracing, and vaccination campaigns have been implemented globally.

The pandemic has also exposed existing vulnerabilities and inequalities within societies, disproportionately affecting marginalized communities, low-income countries, and front-line workers [8]. Disparities in access to healthcare and socioeconomic factors have exacerbated the impact of COVID-19 on vulnerable populations. The rapid development of vaccines, diagnostic tests, therapeutics, and public health interventions has demonstrated the collective resilience and ingenuity of the global scientific community in the face of adversity.

C. Shallow Models' Results

In this section, we will analyze the results from the convex literature for each shallow machine learning model.

K-Nearest Neighbors (KNN)

Ye and colleagues [9] implemented an intelligent system for classifying the severity of COVID-19, to help clinicians

in their decisions. The authors trained the model HHO-FKNN, based on KNN, considering the list of symptoms, complications degree, already existing diseases, and the immune system. They achieved an average accuracy of 94%, the Matthews' correlation coefficient of 88.91%, an average sensitivity of 90%, and an average specificity of 96.67%.

In another article [10], Hamed and coauthors focused on incomplete datasets to predict if a patient suffers from coronavirus or not, and to classify properly, using KNN bases, all the patients. They used two distances: mahalanobis and euclidean. For the mahalanobis distance, the authors obtained an average accuracy of 84%, a sensitivity of 76%, a precision of 95%, and an F1-score of 84%. For the euclidean distance, they achieved an accuracy of 88%, a sensitivity of 87%, a precision of 91%, and the F1-score of 88%.

Decision Trees (DT)

The authors of [11] calculated the performance evaluation metrics for predicting retrospectively the coronavirus considering the blood gas parameter, by using decision trees methods. They got an accuracy ratio of 65% for the correctly predicting cases and 68.2% for correctly identifying people who indeed suffer from coronavirus. When categorizing patients by cutoff values (less than 1.0, between 1.0 and 1.6, and bigger than 1.6), the achieved an accuracy of 92.7%, the metric which was the main target for the authors.

Support Vector Machines (SVM)

In the paper written by Singh and coauthors [12], the scientists tried to predict coronavirus with SVM by treating on time series data. They considered the active cases, total number of deaths, and recovered ones from January and until April 2020 with international data. They referenced to the article [13], where an accuracy of 88% and an F1-score value of 76% were achieved. The authors of [14] got an accuracy of 88.76% by using the radial basis function in SVM when classifying countries into those at risk and without risk.

Classification and Regression Trees (CART)

CART [15] have been utilized in COVID-19 prediction and classification tasks owing to their simplicity and interpretability. By recursively partitioning the data based on the most informative features, CART constructs decision trees that can effectively classify COVID-19 cases into different categories or predict outcomes such as disease severity. CART's ability to handle both numerical and categorical data makes it well-suited for analyzing heterogeneous COVID-19 datasets with diverse epidemiological variables [16].

The authors of [17] built a predictive instruction for COVID-19 pneumonia and classified pneumonia into the one provoked by COVID-19 and not provoked by it. They obtained an area under the ROC curve (ROC-AUC) of 86%, and an accuracy of maximum 95%. On the other hand, Zimmerman and colleagues [16] obtained an ROC-AUC of 76%, a sensitivity of 69%, and specificity of 78%.

Extreme Gradient Boost (XGBoost)

XGBoost [18], an ensemble learning technique, has been widely applied in COVID-19 prediction and classification tasks due to its exceptional performance and scalability. By combining the predictions of multiple weak learners, such as decision trees, XGBoost can effectively capture complex patterns in COVID-19 data and improve predictive accuracy. However, XGBoost may require careful tuning of hyperparameters and regularization techniques to prevent overfitting, especially with large-scale COVID-19 datasets [19].

In the article [20], Carvalho and colleagues built an approach which can diagnose accurately and precisely the COVID-19 for the patients with XGBoost layer added to a convolutional neural network. They obtained an accuracy of 95.07%, a recall of 95.1%, precision of almost 95%, the F1-score and the ROC-AUC of 95% both, while the Cohen's index was 90%. The second article [21] (Fang et al.) related to XGBoost analyzed statistical indicators such as mean squared errors, mean absolute errors and the R-squared coefficient to improve the prediction of the number of patients who are infected with SARS-CoV-2 only in the USA, providing an excellent R-squared, no bigger than 4.1.

D. Ensemble Learning Framework

The predictions of many base models are combined in an ensemble learning process to get a final prediction that is more accurate and dependable [22]. The idea behind ensemble learning is to leverage the diversity of individual models to compensate for their weaknesses and improve overall predictive performance. There are several ensemble learning frameworks, including bagging, boosting, and stacking, each with its advantages and disadvantages.

There are various benefits to bagging. By using diverse subsets of the training data to train numerous base models, it effectively lowers variance and overfitting. Additionally, bagging can be parallelized. However, it comes with its own set of disadvantages. As the number of base models increases, computational complexity and memory requirements may become prohibitive, particularly for large datasets [23]. Boosting, on the other hand, offers distinct advantages. It builds a strong learner iteratively by focusing on examples that are difficult to classify or have high prediction errors. Also, it is sensitive to noise and outliers, potentially leading to overfitting on irrelevant examples during training [24]. Stacking integrates predictions from multiple heterogeneous base models, leveraging the strengths of different modeling techniques. On the other hand, stacking may suffer from information leakage or overfitting if the meta-learner is trained on predictions from the same data used to train the base models [25].

III. PROPOSED MODEL

Given the complexity and heterogeneity of COVID-19 data, a diverse set of shallow models is selected to capture various aspects of the pandemic. Models such as KNN, DT, SVM, CART, and XGBoost (Fig. 2) are chosen based on their

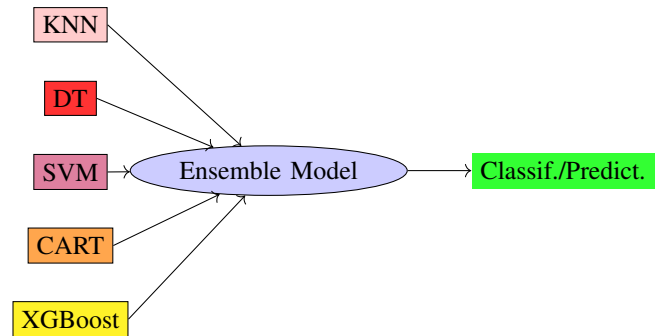


Fig. 2. Hyperparameterized ensemble learning model for classifying countries into those at risk and safe, and predicting the specific prospective *Total Cases* value for a sample.

suitability for handling different types of COVID-19 data, including epidemiological, clinical, genomic, and environmental factors.

Each selected shallow model is trained on COVID-19 data obtained from reliable source publicly available on Kaggle [26]. The training process involves preprocessing the data, selecting appropriate features, and tuning hyperparameters to optimize model performance. For instance, KNN is trained using historical COVID-19 case data to predict future transmission trends, while SVM is trained to classify patients based on clinical symptoms and demographic information. Ensemble learning techniques, including bagging, boosting, and stacking, are employed to combine the predictions of the individual shallow models. Bagging is used to aggregate predictions from multiple models to reduce variance and overfitting, boosting adapts the models iteratively to improve performance over time, and stacking integrates predictions from diverse models to capture complex relationships in COVID-19 data. Hyperparameters for both individual models and the ensemble framework are tuned using COVID-19-specific data and evaluation metrics. Grid search or Bayesian optimization techniques are applied to identify optimal hyperparameter configurations that maximize predictive performance and classification accuracy for COVID-19-related outcomes such as disease transmission, severity, etc.

The performance of the ensemble model is validated using cross-validation techniques on COVID-19 datasets. Special attention is given to account for temporal and geographical variations in COVID-19 data to ensure robustness and generalizability. Evaluation metrics such as accuracy, precision, recall, F1-score, and ROC-AUC are used to assess predictive performance and classification accuracy.

Once validated, the ensemble model is deployed into production systems or applications for real-time COVID-19 prediction and classification (Fig. 2). Integration into public health surveillance systems, decision support tools, or epidemiological models ensures that the ensemble model contributes to informed decision-making and effective public health interventions in the fight against COVID-19.

IV. RESULTS

In this chapter, we will present our results, through which we get a possible solution for the specific total case value for Romania and also for classifying nations into those at risk and without risk. We offer some graphs and charts for results visualization.

The calculations include the statistical indicators: Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared coefficient, and correlation coefficient. We also calculate the evaluation metrics: sensitivity, specificity, accuracy, and precision. For the selected dataset, a strong Pearson correlation is obtained between the column *Total Recovered*, and the column for the dependent variable, *Total Cases* (0.9). Between the column *Active Cases* and *Total Cases*, there is a strong correlation, of 0.72. Other columns included in the independent variable are *Population*, *Serious or Critical*, and *Total Tests*.

A. Statistical Indicators for All Selected Models

We calculate the statistical indicators for 40 different values of the selected parameters of each state-of-the-art model. The *k-neighbors* parameter is the one that varies in the KNN model. The value for MSE does not exceed 2.1 and MAE is at most 1.1. For the 40 values, the R-squared coefficient is between 0.75 and 0.98. The consistently low MSE and MAE values across different *k* values suggest that the KNN model is consistently accurate in its predictions for various levels of *k-neighbors* parameter. This stability in accuracy is important for identifying critical regions exactly. It is important to select an appropriate value, to balance model complexity and generalizability to ensure reliable predictions.

In the DT model, the parameter *max-depth* varies. An MSE between 1 and 4 is obtained; the MAE is not greater than 1.7. The coefficient of determination for DT is between 0.55 and 0.95. The varying R-squared coefficients indicate that certain depths might result in better explanations and stronger relationships in the data, potentially leading to more accurate predictions for critical geographic regions. The choice of an appropriate tree depth involves considering a balance between accuracy and model complexity to ensure reliable predictions.

The SVM model provides, through the regularization parameter $C = \frac{1}{10}$, an MSE of at most 7.25, and an MAE of at most 2.3. The R-squared coefficient is around 0.7 for most values. Although the explanatory strength of the model might be moderate, the consistent accuracy at different levels of regularization indicates that the model performs fairly well. It is important to consider the specific context of the study, the trade-offs between regularization and accuracy, and potential avenues for model improvement.

The CART model, with the value for *max-depth* and *min-leaf* varying, has an MSE between 5 and 11, and an MAE of at most 2.6. The R-squared coefficient ranges between 0.88 and 0.98. It is important to select appropriate combinations of *max-depth* and *min-leaf* for dependable forecasts, to strike a balance between generalizability and model complexity.




Fig. 3. The MSE indicator for the *Renesansa* ensemble model.




Fig. 4. Mean Absolute Error (MAE) for *Renesansa* ensemble model.

XGBoost is an outstanding model, with an MSE no higher than 1.8 and an MAE between 0.3 and 1. The coefficient of determination is between 0.82 and 0.98, and the correlation coefficient is between 0.91 and 0.99.

When it comes to important assessment measures such as MSE (Fig. 3), MAE (Fig. 4), and R-squared, *Renesansa* performs admirably. *Renesansa* has an exceptionally low MSE (between 0.001 and 1.19), indicating small squared discrepancies between expected and actual values. This implies that the model offers extremely precise evaluations of the risk of COVID-19 in a nation, which is essential to inform public health initiatives and policy choices. In a similar vein, the model shows a low MAE, focusing on small changes between the predicted values and the actual values. This demonstrates how accurately *Renesansa* can estimate the risk variables for COVID-19, allowing policymakers and health authorities to make well-informed decisions. *Renesansa* also produces a high R-squared value (between 0.975 and 0.9997), suggesting that the model explains a substantial amount of the variability in




Fig. 5. Worldometer [27] value for Romania on 18th April 2023.



Fig. 6. Total cases value obtained by model.

COVID-19 risk variables.

This highlights the model’s capacity to identify underlying trends and patterns in COVID-19 data, which is crucial for formulating workable plans to stop the virus’s spread and handle public health emergencies. Renesansa’s notable performance in evaluating the risk of COVID-19 for nations is highlighted by its remarkable results in MSE, MAE, and R-squared. While a high R-squared value indicates a great power of explanations, low MSE and MAE values indicate accurate predictions and minimum mistakes. These findings highlight the significance of sophisticated ensemble models such as Renesansa in directing evidence-based global responses to the COVID-19 epidemic.

The *Total Cases* value for Romania on 18 April 2023 on Worldometer [27] was 3,382,756. We can observe the result for Romania, obtained by model’s prediction, which is very

similar to the one from Worldometer (3,382,872). There is an extremely high performance of this model in the set of independent variable columns (compare the results of Fig. 5 and Fig. 6).

B. Binary Classification Considering the Proportion Between Total Cases and Population

The results obtained for the performance evaluation metrics of the models are strong (Fig. I). Performance evaluation metrics provide insight into the strengths and weaknesses of each model in the context of detecting critical geographic regions in the COVID-19 dataset. These metrics provide a comprehensive view of each model’s performance. The best model for a specific use case might depend on the priorities: whether it is achieving high accuracy, minimizing false positives, maximizing sensitivity, or maintaining a balance between different metrics.

TABLE I
PERFORMANCE EVALUATION METRICS FOR MODELS IN TERMS OF COUNTRY’S COVID-19 RISK CLASSIFICATION.

	Sensitivity	Specificity	Accuracy	Precision
KNN	0.9286	0.9741	0.9606	0.9381
DT	0.8730	0.8427	0.8485	0.5670
SVM	0.9950	0.7614	0.7788	0.2474
CART	0.8846	0.8889	0.8879	0.7113
XGBoost	0.9970	0.9549	0.9667	0.8866
Renesansa	0.9975	0.9749	0.9818	0.9381

Important information is revealed by the evaluation results (see Table I) of several machine learning models used to group nations into COVID-19 risk categories. Countries, whose ratio between *Total Cases* and *Population* is above 0.003 (0.3%), are seen as areas at risk of infection. This value is calculated with Youden J Index technique. To determine a cutoff using the Youden’s J Index, sensitivity and specificity of the diagnostic test are initially assessed. Subsequently, Youden’s J Index is computed by summing the sensitivity and specificity, then subtracting one. Ultimately, the threshold value that yields the highest Youden’s J Index is chosen as the cutoff, signifying the most favorable equilibrium between sensitivity and specificity in the diagnostic evaluation [28]. All metrics show that the ensemble learning model has the highest performance (99.75% sensitivity, 97.49% specificity, 98.18% accuracy, and 93.81% precision) demonstrating its reliability in differentiating between nations that are at risk and those that are not. This shows that the accuracy and dependability of the predictions can be improved by integrating various models. Notably, XGBoost also performs admirably, especially when it comes to sensitivity and specificity (99.7% and 95.49%), which are crucial for accurately identifying nations that are actually at danger while reducing false positives. SVM and decision tree models, on the other hand, show less accuracy (76.14% and 84.27%), suggesting a higher false alarm rate. In order to properly identify at-risk and non-at-risk nations, sensitivity and specificity are essential. These findings highlight how crucial trustworthy prediction models are in directing appropriate actions.

V. CONCLUSIONS

We find interesting paths to enhance classification accuracy and prediction performance by using ensemble learning approaches for COVID-19 prediction and classification tasks. Through the combination of a wide range of shallow models, such as DT, SVM, CART, KNN, and XGBoost, we have shown the ability to improve performance and mitigate the shortcomings of individual models on a variety of COVID-19-related datasets.

Our results highlight how crucial ensemble learning frameworks—like bagging, boosting, and stacking—are for efficiently combining predictions from several models to identify the intricate patterns and correlations present in COVID-19 data. We have demonstrated by thorough hyperparameter tuning, model validation, and interpretation analysis that ensemble learning models provide reliable solutions for this topic.

Our technique makes it easier to comprehend and evaluate model outputs by offering insights into the variables influencing COVID-19 predictions and classifications. We may improve real-time tracking, forecasting, and reaction efforts in the ongoing fight against the COVID-19 pandemic by integrating the ensemble model into decision support systems.

ACKNOWLEDGMENT

The author thanks the Computer Science Department of the West University of Timisoara for the support in terms of resources and some professors for the indicated suggestions.

REFERENCES

- [1] Ameer Sardar Kwekha-Rashid, Heamn N Abduljabbar, and Bilal Al-hayani. Coronavirus disease (covid-19) cases analysis using machine-learning applications. *Applied Nanoscience*, 13(3), 2023. DOI: 10.1007/s13204-021-01868-7.
- [2] Hafsa Barea Syeda, Mahanazuddin Syed, Kevin Wayne Sexton, Shorabuddin Syed, Salma Begum, Farhanuddin Syed, Fred Prior, and Feliciano Yu Jr. Role of machine learning techniques to tackle the covid-19 crisis: systematic review. *JMIR medical informatics*, 9(1):e23811, 2021. DOI: 10.2196/23811.
- [3] Sara Platto, Tongtong Xue, and Ernesto Carafoli. Covid19: an announced pandemic. *Cell Death & Disease*, 11(9):799, 2020. DOI: 10.1038/s41419-020-02995-9.
- [4] Mustafa Hasöksüz, Selcuk Kilic, and Fahriye Saraç. Coronaviruses and sars-cov-2. *Turkish journal of medical sciences*, 50(9):549–556, 2020. DOI: 10.3906/sag-2004-127.
- [5] World Health Organization et al. Coronavirus disease 2019 (covid-19): situation report, 116. 2020. DOI: 10.2139/ssrn.3566298.
- [6] Marco Ciotti, Massimo Ciccozzi, Alessandro Terrinoni, Wen-Can Jiang, Cheng-Bin Wang, and Sergio Bernardini. The covid-19 pandemic. *Critical reviews in clinical laboratory sciences*, 57(6):365–388, 2020. DOI: 10.1080/10408363.2020.1783198.
- [7] Rakesh Padhan and KP Prabheesh. The economics of covid-19 pandemic: A survey. *Economic analysis and policy*, 70:220–237, 2021. DOI: 10.1016/j.eap.2021.02.012.
- [8] Walter Cullen, Gautam Gulati, and Brendan D Kelly. Mental health in the covid-19 pandemic. *QJM: An International Journal of Medicine*, 113(5):311–312, 2020. DOI: 10.1093/qjmed/hcaa110.
- [9] Hua Ye, Peiliang Wu, Tianru Zhu, Zhongxiang Xiao, Xie Zhang, Long Zheng, Rongwei Zheng, Yangjie Sun, Weilong Zhou, Qinlei Fu, et al. Diagnosing coronavirus disease 2019 (covid-19): Efficient harris hawks-inspired fuzzy k-nearest neighbor prediction methods. *IEEE Access*, 9:17787–17802, 2021. DOI: 10.1109/access.2021.3052835.
- [10] Ahmed Hamed, Ahmed Sobhy, and Hamed Nassar. Accurate classification of covid-19 based on incomplete heterogeneous data using a kn variant algorithm. *Arabian Journal for Science and Engineering*, 46:8261–8272, 2021. DOI: 10.1007/s13369-020-05212-z.
- [11] Mehmet Tahir Huyut and Hilal Üstündağ. Prediction of diagnosis and prognosis of covid-19 disease by blood gas parameters using decision trees machine learning model: a retrospective observational study. *Medical gas research*, 12(2):60–66, 2022. DOI: 10.4103/2045-9912.326002.
- [12] Vijander Singh, Ramesh Chandra Poonia, Sandeep Kumar, Pranav Dass, Pankaj Agarwal, Vaibhav Bhatnagar, and Linesh Raja. Prediction of covid-19 corona virus pandemic based on time series data using support vector machine. *Journal of Discrete Mathematical Sciences and Cryptography*, 23(8):1583–1597, 2020. DOI: 10.1080/09720529.2020.1784535.
- [13] Y Lebrini, A Boudhar, R Hadria, H Lionboui, L Elmansouri, R Arrach, P Ceccato, and T Benabdelouahab. Identifying agricultural systems using svm classification approach based on phenological metrics in a semi-arid region of morocco. *Earth Systems and Environment*, 3(2):277–288, 2019. DOI: 10.1007/s41748-019-00106-z.
- [14] Sajja Tulasi Krishna and Hemantha Kumar Kalluri. Lung image classification to identify abnormal cells using radial basis kernel function of svm. In *Smart Technologies in Data Science and Communication: Proceedings of SMART-DSC 2019*, pages 279–285. Springer, 2020. DOI: 10.1007/978-981-15-2407-333.
- [15] Leo Breiman, Jerome Friedman, Richard Olshen, and Charles Stone. *Cart. Classification and regression trees*, 1984. DOI: 10.1201/9781315139470-8.
- [16] Richard K Zimmerman, Mary Patricia Nowalk, Todd Bear, Rachel Taber, Karen S Clarke, Theresa M Sax, Heather Eng, Lloyd G Clarke, and GK Balasubramani. Proposed clinical indicators for efficient screening and testing for covid-19 infection using classification and regression trees (cart) analysis. *Human Vaccines & Immunotherapeutics*, 17(4):1109–1112, 2021. DOI: 10.1080/21645515.2020.1822135.
- [17] Sayato Fukui, Akihiro Inui, Takayuki Komatsu, Kanako Ogura, Yutaka Ozaki, Manabu Sugita, Mizue Saita, Daiki Kobayashi, and Toshio Naito. A predictive rule for covid-19 pneumonia among covid-19 patients: A classification and regression tree (cart) analysis model. *Cureus*, 15(9), 2023. DOI: 10.7759/cureus.45199.
- [18] Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, Rory Mitchell, Ignacio Cano, Tianyi Zhou, et al. Xgboost: extreme gradient boosting. *R package version 0.4-2*, 1(4):1–4, 2015. DOI: 10.32614/cran.package.xgboost.
- [19] Junling Luo, Zhongliang Zhang, Yao Fu, and Feng Rao. Time series prediction of covid-19 transmission in america using lstm and xgboost algorithms. *Results in Physics*, 27:104462, 2021. DOI: 10.1016/j.rinp.2021.104462.
- [20] Edelson Damasceno Carvalho, Edson Damasceno Carvalho, Antonio Oseas de Carvalho Filho, Flávio Henrique Duarte de Araújo, and Ricardo de Andrade Lira Rabêlo. Diagnosis of covid-19 in ct image using cnn and xgboost. In *2020 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–6. IEEE, 2020. DOI: 10.1109/iscc50000.2020.9219726.
- [21] Zheng-gang Fang, Shu-qin Yang, Cai-xia Lv, Shu-yi An, and Wei Wu. Application of a data-driven xgboost model for the prediction of covid-19 in the usa: a time-series study. *BMJ open*, 12(7):e056685, 2022. DOI: 10.1136/bmjopen-2021-056685.
- [22] Thomas G Dietterich et al. Ensemble learning. *The handbook of brain theory and neural networks*, 2(1):110–125, 2002. DOI: 10.7551/mitpress/3413.001.0001.
- [23] Leo Breiman. Bagging predictors. *Machine learning*, 24:123–140, 1996. DOI: 10.1007/bf00058655.
- [24] Robert E Schapire et al. A brief introduction to boosting. In *Ijcai*, volume 99, pages 1401–1406. Citeseer, 1999. DOI: 10.1007/3-540-49097-31.
- [25] Kai Ming Ting and Ian H Witten. Stacking bagged and dagged models. 1997. DOI: 10.1109/icdm.2010.49.
- [26] Covid-19 dataset. <https://www.kaggle.com/datasets/selfishgene/covid19-worldometer-snapshots-since-april-18?resource=download>, last accessed on 18 mar.
- [27] Worldometer information about coronavirus. <https://www.worldometers.info/coronavirus/>, last accessed on 17 mar.
- [28] Ronen Fluss, David Faraggi, and Benjamin Reiser. Estimation of the youden index and its associated cutoff point. *Biometrical Journal: Journal of Mathematical Methods in Biosciences*, 47(4):458–472, 2005. DOI: 10.1002/bimj.200410135.

Agricultural Data Space: the METRIQA Platform and a Case Study in the CODECS project

Manlio Bacco

*Inst. of Information Science and Technologies (ISTI)
National Research Council (CNR), Pisa, Italy
0000-0001-6733-1873*

Alexander Kocian

*Dept. Computer Science
University of Pisa, Italy
0000-0001-8847-0768*

Antonino Crivello

*Inst. of Information Science and Technologies (ISTI)
National Research Council (CNR), Pisa, Italy
0000-0001-7238-2181*

Marco Gori

*Dept. of Information Engineering and Mathematics (DIISM)
University of Siena, Italy
0000-0001-6337-5430*

Giovanna Maria Dimitri

*Dept. of Information Engineering and Mathematics (DIISM)
University of Siena, Italy
0000-0002-2728-4272*

Paolo Barsocchi

*Inst. of Information Science and Technologies (ISTI)
National Research Council (CNR), Pisa, Italy
0000-0002-6862-7593*

Gianluca Brunori

*Dept. of Agriculture, Food and Environment
University of Pisa, Italy
0000-0003-2905-9738*

Stefano Chessa

*Dept. Computer Science, University of Pisa, Italy
ISTI-CNR, Italy
0000-0002-1248-9478*

Abstract—This work describes the ongoing design and development of the METRIQA platform, hosting the Italian agrifood data space. Both are key components that the Italian National Research Centre for Agricultural Technologies is putting forward in its activities. We present a high-level description of the platform, which is designed to provide web-like access to digital resources and services following an approach called Web of Agri-Food, to support the digital transformation of the sector in Italy. To show its potential, we also present a real case study demonstrating both the benefits and impacts of the proposed architecture, connecting stakeholders and authorities at different levels.

Index Terms—agri-food, data space, agriculture, traceability, data sharing

I. INTRODUCTION

THE European Green Deal is part of the European Union (EU)'s response to the Sustainable Development Goals (SDGs), and it sets specific targets for the agricultural sector, among others, emphasizing the potential for the adoption of digital technologies. Member States, like Italy, are responsible for implementing these objectives, funded by the National Recovery and Resilience Plan (PNRR) and being implemented

This work has been funded within the EU Horizon Europe research and innovation programme under GA no. 101060179 and the project AGRITECH Spoke 9 - Codice progetto MUR: AGRITECH "National Research Centre for Agricultural Technologies" - CUP CN00000022, of the National Recovery and Resilience Plan (PNRR) financed by the European Union "Next Generation EU".

through the National Research Centre for Agricultural Technologies (AGRITECH). The latter is organized into nine thematic areas called "Spokes," with Spoke 9 focusing on "new technologies and methodologies for traceability, quality, safety, measurements, and certifications to enhance the value and protect the typical traits in agri-food chains." In the activities of spoke 9, the METRIQA technological platform [1] emerges as a solution to support the digital transformation of Italy's agri-food sector. The primary objective of METRIQA is to implement the "Web of Agri-Food" (WoA), which will provide web-like access to digital resources and services. The WoA integrates several data sources, connecting their resources regardless of internal technology and structure. This allows various providers to make unstructured and heterogeneous data available to users, such as the actors in the agri-food sector and other stakeholders. METRIQA users will rely on novel AI-based retrieval services to obtain relevant information, benefiting both research and sectoral stakeholders. METRIQA will also assist private companies in sharing digital resources. A core component of the METRIQA platform is a *data space* [2], which will be used at the national level for traceability and certification services in the test phase of the system.

The EU data strategy [3] aims at fostering the creation and use of data spaces for data integration. Data spaces, which can be described as decentralised data ecosystems focusing on data relationships rather than data management, have the potential to play a crucial role in facilitating a cooperative environment.




Fig. 1: Conceptual description of Data Spaces: data, infrastructure, and services shared for use according to predefined policies in a common data market.

Such ecosystems rely on mutually agreed software building blocks. The vision is that data remains at the source, and using *connectors* (specialized software components) [4] they can be accessed by interested parties (data consumers). Figure 1 shows a conceptual representation of the main functionalities offered by data spaces. Rules for data access and use are established by data owners in the form of policies. A central authority is in charge of the identity verification of connected entities and other functionalities, which will be further discussed in the following of this work.

Notable initiatives on the road to data spaces are worth mentioning, such as the IDS (International Data Spaces), the IDSA (International Data Spaces Association), and DSBA (Data Space Business Alliance) that puts together the GAIA-X European Association for Data and Cloud AISBL, BDVA, FIWARE, and IDSA. There are initiatives also outside the EU, such as the Data Society Alliance (DSA) in Japan. We remark the the DSSC (Data Spaces Support Centre) initiative, and the development of a smart platform at the EU level, namely SIMPL [5] to further support the development, testing, and use of data spaces through the development of a common middleware solution. The idea is to provide a reference implementation to facilitate the exchange of data across company boundaries and temporarily address challenges such as interoperability, transparency, trust, and data security. Coupled with the use of W3C semantic standards, data spaces aim at formal and machine-interpretable specifications of concepts for unified data understanding. The use of e.g., MyData Control Technologies [6], which provides a technical implementation of data sovereignty, the aforementioned control on data by owners can be exercised.

In this work, we explore the current state of play and ongoing initiatives when it comes to data spaces, especially in the agricultural sector, in Section II, and provide details on the ongoing efforts for the development of the aforementioned METRIQA platform in Section III. Additionally, we explore a real case study, located in Tuscany, Italy, in Section IV and discuss how data spaces can provide benefits in such a context. We conclude the work in Section V.

II. STATE OF THE ART

Digital information is a key driver of business innovation, social innovation and prosperity, thus it plays an important

role in transforming our lives. One of the driving forces behind this transformation is the exponential growth of data in recent years, considering the wide spread of the Internet of Things devices and, in general, low-cost and commercial hardware widely spread and daily used. Digitalization has also arrived in the agriculture sector [7]. As an example, high-tech greenhouses have become versatile and multifunctional environments equipped with sensors that acquire data related to the microclimatic environment as well as the status of the plant. A specially designed "computer brain" learns from the sensed data to infer the health status of the crops statistically [8], and to perform (optimum) control automatically [9]. More advanced closed-loop high-tech greenhouses have moved information processing to a dedicated cloud but the infrastructure is scattered, unstructured and vertical [10]. To improve efficiency, current high-tech solutions need to be capable of inter-operating and communicating with other systems and clouds making an infrastructure-based approach essential. As a consequence, data platforms are creating a marketplace able to connect data providers and data users.

Focusing on data space technologies, semantic World Wide Web Consortium (W3C) standards, such as the resource description framework (RDF) and Web Ontology Language (OWL) have been developed over two decades, to ensure a common understanding of the data shared in the data space. The IDS information model is based on a graph-based RDF ontology, to link the conceptual functionalities, roles, and processes with the implementations in the Connector interfaces and endpoints. MyData Control Technologies can be applied to implement the capability of data sovereignty. To share data simultaneously, there must be some form of usage control enforcement. This concept relies on a certification process for components and environments based on public key infrastructure. The emergence of federated digital ecosystems has the potential to efficiently improve data availability and also offer new strategies in dealing with large volumes of data, as required by model training for AI-based approaches. Anyway, a key hurdle on the road of such an ecosystem is the availability of trustworthy data sharing and management [11]. As argued in [12], trust should be defined as a function of time, encompassing a lengthy, intricate, and continuously adapted chain of trust between services, providers, and users to adapt to the high dynamicity of the data market.

Several studies (e.g., [13]–[16]) and EU-funded projects have studied and demonstrated new approaches and solutions for digitalising agriculture in recent years, and we cite a few of them in the following. For instance, the H2020 DESIRA project (GA 818194) has explored the role of digital technologies in agriculture, forestry, and rural areas, highlighting both the potential and the risks linked with their increasing use [17]. The HE Data4Food2030 (GA 101059473) project focuses on the data economy and the HE CODECS project (GA 101060179) is carrying out a cost-benefit analysis of technologies in the agricultural sector, assessing changes in the socio-economic process by comparing the situation before and after the introduction of digital technologies in 20 case studies all over Europe [18]. The “OPEN DEI” initiative presents a 6C (connection, cyber, computing, content/context, community, customization) approach in the form of a reference architecture framework designed to facilitate cross-domain digital transformation [19]. Such a framework builds on the extensive use of sensing units in most scenarios nowadays, and the need to have data-driven pipelines and workflow management to fully exploit the huge amount of available data for decision support, especially in the industrial sector. With specific reference to agriculture, it is worth mentioning the ongoing activities of the HE AgriDataSpace action (GA 101083401) and of the HE AgriDataValue project (GA 101086461), which consider data spaces as a key infrastructural component for secure and trusted data exchanges in the EU. At the national level, we point interested readers to valuable initiatives, such as SIEX¹ in Spain, a set of interconnected databases and administrative registers, with information about agrarian holdings, and to the GC4SHEEP² initiative when it comes to livestock.

III. ITALIAN AGRIFOOD DATA SPACE (IADS)

The National Research Centre for Agricultural Technologies (AGRITECH) in Italy aims at developing “new technologies and methodologies for traceability, quality, safety, measurements and certifications to enhance the value and protect the typical traits in agri-food chains”. The centre aims to “create an information platform integrating all work packages”. Following this approach, the technological platform METRIQA (MEasurements, TRaceability and Quality in Agri-food chains) is designed to offer a wide spectrum of services for research, agri-food companies, citizens and public bodies.

In particular, from the point of view of services to research, METRIQA provides cloud storage to keep the data and information produced by the research activity of the Agritech Centre and makes use of Artificial Intelligence (AI) to enrich such digital resources with metadata [20], to discover knowledge, and to construct application-specific decision support systems, sharing similarities with the web but tailored for the agri-food sector. Furthermore, it also offers personalized notebooks allowing for the advanced analysis of research data to build

novel AI models and tools. At the same time, METRIQA implements the abstraction we named WoA to enable other stakeholders to access the information using search engines and conversational engines based on GPT (Generative Pre-training Transformer) technologies. As an example of the services provided to research, consider the scenario of a research centre that is operating a study for the certification of the origin of a specific product, by leveraging a data collection campaign on fields to collect samples of soil and biological material from the place of production, and on the analysis of these samples to build a hyper-spectral profile of the products of the area. The results of this research and the methodology to prove the origin are published in a technical report for other experts. METRIQA can support this research by providing a place where to store all this information (reports and datasets), and provide the tools to process such data with advanced AI tools. At the same time however, the ability of METRIQA to index the report and to embed its knowledge into a conversational chatbot allows even non-researchers to access the information about the certification of origin of that product. In turn, this information may be used by companies to provide guarantees about the origin of their products to their consumers.

From the point of view of services to agri-food companies, METRIQA offers two main services: the support to traceability of products/goods leveraging on blockchain technologies and the design and implementation of the IADS according to the specifications of IDS mentioned above. The objective is to provide a scalable data space that can grow over time with the contributions of individual project partners as well as service providers. To this purpose, METRIQA encompasses several nodes, all run by the Agritech centre, such as an authority node (that serves also as an identity provider), a broker, a transaction node, and several others interconnected through suitable connectors. Among those, a *public node* will offer services to research, citizens and public bodies as discussed above. Four classes of stakeholders have been identified, namely service and data producers, consumers, node managers, and the governance body (the Agritech Centre). Figure 2 shows a high-level view of METRIQA, highlighting that the data space is a key component of the platform. Each participant must agree with the Data Space policy at registration. The registration process is controlled by the (trusted) authority node.

The other nodes can be servers or clusters of servers located in the cloud or at the network edge, managed by other public or private entities. Each of these nodes hosts data and services specific to the needs of the respective participants. In particular, each of these nodes may implement additional services tailored to companies or specific production processes. For example, through their membership to the data space, companies along a production and distribution chain can share traceability data, enabling food traceability services from grower to consumer. Integration of data from different sources and export through various services is required for such services. Service and resource providers populate the platform with data, knowledge, and services, including re-

¹Agrarian Holding Information System: www.fega.gob.es/en/content/sieux

²Federated Data Cloud Platform with Artificial Intelligence Layer for the Genetic and Reproductive Improvement of the National Dairy Sheep at <https://gc4sheep.com>



Fig. 2: A high-level visualisation of the METRIQA platform and its nodes.

searchers, companies, professionals, and public bodies. End users consume data, information, and services offered by the platform. Node managers are entities running nodes, which can include IT companies, individual researchers, and the Agritech Centre itself.

At the current stage of the project, we have designed the reference architecture of METRIQA, specifically the three layers of storage, components, and applications. Starting with the Storage Layer, its main components are the message broker and the stored data. The storage contains the data lake of the (research) data that Spoke 9 is generating. At the Components Layer, METRIQA implements a decision support system offering web-style retrieval services (such as a natural language question-answering engine), modelling notebooks using interpreted languages (e.g. Python) for the exploratory data analysis, APIs for the data flow among data providers and data consumers, metadata enrichment, Software as a Service (SaaS) and Platform as a Service (Paas) models, and containers of sectorial decision support systems internally developed. Presently, the development is focusing on the public node, using the cloud services of the Italian provider "Consorzio Interuniversitario dell'Italia Nord Est per il Calcolo Automatico" (CINECA). The development of the public node includes the implementation of a (limited) number of decision support systems aimed at supporting the research activities of the centre, and the deployment of the main elements that support the creation of the data space.

IV. REAL CASE STUDY IN TUSCANY, ITALY

The case study described in what follows comes from the ongoing activities in one of the living labs³ of the EU project CODECS cited above. The project puts forward a vision of sustainable digitalisation in the field of agriculture to carefully consider the existing barriers and the potential negative impacts [17] which the push of digital technologies may trigger and builds on the results of the DESIRA project, such as co-design activities to reduce risks and maximise benefits for farmers and other stakeholders [21].

A representation of the case study is provided in Fig. 3 from the viewpoint of data spaces. From the left, actors in the field (farmers, farmers' associations, factories, and so on) are shown: they represent an ecosystem of milk producers and cheese makers in southern Tuscany. Farming requires lots of paperwork for compliance with regulations, and the provision of data to control entities is a crucial activity to have access to incentives, such as those foreseen in the EU CAP (Common Agricultural Policy). Most of the data and information are not natively digital yet, meaning that several paper documents are compiled and exchanged among actors and public bodies. An added complexity is because similar data are often requested by more than one control entity, which translates into being forced to provide the same information in different formats to different recipients through different procedures. Such a documentation burden represents a significant hurdle, especially

³Living Labs are networks of farmers, knowledge intermediaries, stakeholders, and policymakers to address agricultural challenges with a system-level approach to provide insights to policymakers and support sustainable farming practices.

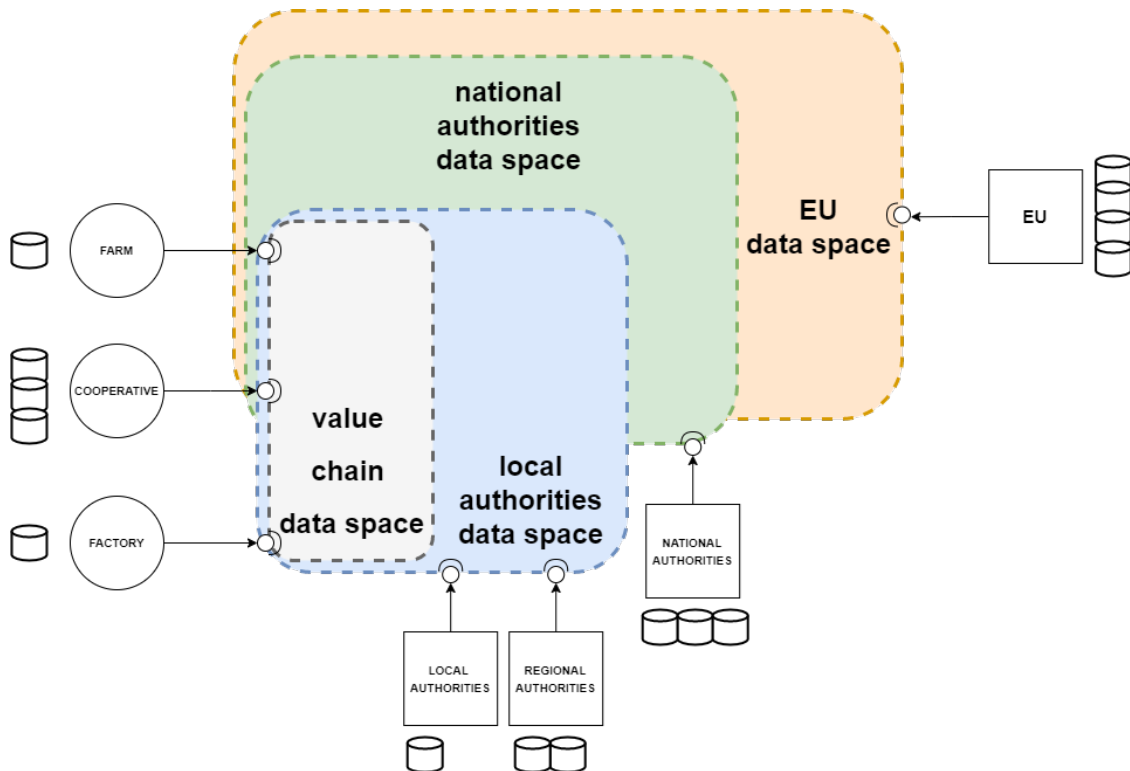


Fig. 3: A schematisation of the case study under consideration, highlighting several data spaces at different levels, nested and overlapped to meet different needs.

for small farmers. Farm Management and Information Systems (FMISs) can be of strong support in this regard, but not all portals and platforms for declarations and certifications can be programmatically queried nor automatic upload procedures can be carried out. Again, this means that most procedures must be carried out manually. Looking again at Fig. 3, what just presented translates into authorities (at the bottom) not being *connected* to the farms or to the advisory services farms rely upon. In other words, the lack of a programming interface (in the form of a standard API) may have severe impacts on how burdening the work can be. An additional *connection* must be foreseen between national and EU authorities (i.e., between the national agency for agricultural subsidies and the EU CAP system), although on a different scale.

In such a scenario, the use of one (or multiple, as in Fig. 3) data space may represent a viable solution to *connect* the aforementioned systems and actors by using standard and interoperable software components (connectors) that support data sharing and data exchange functionalities. Going back to the proposed case study, it means that FMISs at the farm level have connectors towards e.g. a value chain (or local) data space enabling data sharing among milk producers, the cheese factory, the farmers' cooperative, and other relevant actors (e.g., an external laboratory for milk testing). A local space can ensure that data are willingly shared by each actor according to predefined rules, thus feeding and being fed by the different management systems in use. A part of those data is of interest

to local authorities (for instance, the Tuscany region in this case) in charge of e.g., control and certification procedures, thus an additional data space could be set up for data sharing and exchange, with different rules set by data owners to limit the access by the regional authorities to data as required by existing regulations in the area. Data spaces can be nested and overlapped, as depicted in Fig. 3 and discussed in [22], two features that should be strongly exploited in our opinion, as we hinted above. In such a way, data can be shared across multiple groups of actors according to different requirements, rules, and needs by setting up a single connector e.g. at the farm level instead of setting up multiple connectors to connect to different data spaces or having a single data space spanning from the local to the EU level. In Fig. 3, the overlapping and nesting of data spaces are represented by different colours and partial overlapping of rounded areas, each representing a different space.

V. CONCLUSIONS

The METRIQA platform is under active development within the activities of PNRR in Italy. The potential of data is being explored in both the literature and the market, and we reported a real case study in which data spaces are supposed to provide benefits. In the foreseeable future, we plan to connect the information systems at the case study level to the METRIQA platform, thus integrating data and putting *connectors* in place.

We expect two broad families of benefits. On the one hand, such convergence will allow the actors on the ground (e.g., in the case study) to interconnect with other companies nationwide, with the perspective of connecting in a (semi-)automatic fashion to additional supply chains. On the other hand, said convergence will open to deeper integration with the centre research services, as the data produced by national companies will be made available to research and industry according to rules the data owners will set, in order to experiment with innovative, AI-based data analysis methodologies.

REFERENCES

- [1] S. Chessa, G. M. Dimitri, M. Gori, and A. Kocian, "WoA: an Infrastructural, Web-based Approach to Digital Agriculture," in *Proc. 14th Int. Symposium on Ambient Intelligence*. Guimarães, Portugal: Springer, 2023, p. 10.
- [2] M. Atzori, A. Ciaramella, C. Diamantini, B. Martino, S. Distefano, T. Facchinetti, F. Montecchiani, A. Nocera, G. Ruffo, R. Trasarti *et al.*, "Dataspaces: Concepts, Architectures and Initiatives," in *CEUR WORKSHOP PROCEEDINGS*, vol. 3606. CEUR-WS, 2024.
- [3] "European Data Strategy," 2019. [Online]. Available: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-data-strategy_en
- [4] J. Pampus, B.-F. Jahnke, and R. Quensel, "Evolving Data Space Technologies: Lessons Learned from an IDS Connector Reference Implementation," in *Leveraging Applications of Formal Methods, Verification and Validation. Practice: 11th International Symposium, ISOFA 2022, Rhodes, Greece, October 22–30, 2022, Proceedings, Part IV*. Springer, 2022, pp. 366–381.
- [5] EU, "Simpl: Cloud-to-Edge Federations Empowering EU Data Spaces," Tech. Rep., 2024. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/policies/simpl>
- [6] L. Nagel and D. Lycklama, "How to Build, Run, and Govern Data Spaces," in *Designing Data Spaces: The Ecosystem Approach to Competitive Advantage*. Springer International Publishing Cham, 2022, pp. 17–28.
- [7] M. Bacco, P. Barsocchi, E. Ferro, A. Gotta, and M. Ruggeri, "The Digitisation of Agriculture: a Survey of Research Activities on Smart Farming," *Array*, vol. 3, p. 100009, 2019.
- [8] A. Kocian and L. Incrocci, "Learning from Data to Optimize Control in Precision Farming," *Stats*, vol. 3, pp. 239–245, 2020, editorial.
- [9] G. Burchi, S. Chessa, F. Gambineri, A. Kocian, D. Massa, P. Milano, P. Milazzo, L. Rimediotti, and A. Ruggeri, "Information Technology Controlled Greenhouse: A System Architecture," in *Proc. IoT Vertical and Topical Summit for Agriculture*. Tuscany, Italy: IEEE, May 2018.
- [10] A. Kocian, G. Carmassi, F. Cela, S. Chessa, P. Milazzo, and L. Incrocci, "IoT-based Dynamic Bayesian Prediction of Crop Evapotranspiration in Soilless Cultivations," *Computer and Electronics in Agriculture*, vol. 205, Feb. 2023.
- [11] M. Koch, S. Kober, S. Straburzynski, B. Gaunitz, and B. Franczyk, "Federated Learning for Data Trust in Logistics," in *FedCSIS (Position Papers)*, 2023, pp. 51–58.
- [12] C. A. Ardagna, N. Bena, N. Bennani, C. Ghedira-Guegan, N. Grecchi, and G. Vargas-Solar, "Revisiting Trust Management in the Data Economy: A Roadmap," *IEEE Internet Computing*, no. 01, pp. 1–8, 2024.
- [13] J. Doerr, R. Kalmar, B. Rauch, and S. Stiene, "Data Spaces in Agriculture - Status Quo and Perspectives," in *LAND.TECHNIK 2022*. VDI Verlag, 2022, pp. 511–520.
- [14] R. Kalmar, B. Rauch, J. Dörr, and P. Liggesmeyer, *Agricultural Data Space*. Springer, 2022, ch. Designing Data Spaces.
- [15] R. Falcão, R. Matar, B. Rauch, F. Elberzhager, and M. Koch, "A Reference Architecture for Enabling Interoperability and Data Sovereignty in the Agricultural Data Space," *Information*, vol. 14, no. 3, p. 197, mar 2023.
- [16] M. Šestak and D. Copot, "Towards Trusted Data Sharing and Exchange in Agro-Food Supply Chains: Design Principles for Agricultural Data Spaces," *Sustainability*, vol. 15, no. 18, p. 13746, 2023.
- [17] A. Ferrari, M. Bacco, K. Gaber, A. Jedlitschka, S. Hess, J. Kaipainen, P. Koltsida, E. Toli, and G. Brunori, "Drivers, Barriers and Impacts of Digitalisation in Rural Areas from the Viewpoint of Experts," *Information and Software Technology*, vol. 145, p. 106816, 2022.
- [18] C. Mannari, M. Bacco, A. Ferrari, L. Ortolani, M. B. Lai, C. Mignani, A. Silvi, A. Malizia, and G. Brunori, "A Methodology for Process Modelling in Living Labs to Foster Agricultural Digitalisation," in *2023 IEEE International Workshop on Metrology for Agriculture and Forestry (MetroAgriFor)*. IEEE, 2023, pp. 19–24.
- [19] A. Kung, S. Gusmeroli, G. Monteleone, A. Dognini, L. Nicolas, and C. Polcaro, "Reference Architectures and Interoperability in Digital Platforms," H2020 Open Dei, techreport, Apr. 2022.
- [20] M. Bauer and C. Augenstein, "Can Unlabelled Data Improve AI Applications? A Comparative Study on Self-Supervised Learning in Computer Vision," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 93–101.
- [21] M. Bacco, A. Ferrari, and G. Brunori, "Co-design of Technological Solutions for Agriculture and Rural areas: Methodology and Cases for Responsible Innovation," in *IEEE 8th World Forum on Internet of Things - Vertical Track: Agriculture, Yokohama, Japan, Nov 2022*.
- [22] L. Nagel, D. Lycklama, and U. Ahle, "Design Principles for Data Spaces: Position Paper," *International Data Spaces Association*, 2021.

Reinforcement Learning based Intelligent System for Personalized Exam Schedule

Marco Barone
Università Giustino Fortunato
Università degli Studi di Foggia
Email: marco.barone@unifg.it

Matteo Ciaschi
National Research Council (CNR)
email: matteo.ciaschi@cnr.it
ORCID: 0009-0009-5119-3563

Zaib Ullah
Università Giustino Fortunato
email: z.ullah@unifortunato.eu

Armando Piccardi
Università Giustino Fortunato
email: a.piccardi@unifortunato.eu

Abstract—Personalized learning has been proving to be useful concept in the learning of a student. Artificial Intelligence (AI) which has revolutionized many aspects of our lives has also been glowingly used in the education sector. One of the fascinating AI technique, the Reinforcement Learning (RL) is considered as the perfect tool to develop personalized solution in the education. RL algorithms have the ability to take into account personal characteristics of each student. This work presents the development of personalized exam scheduler using RL. The intelligent examination scheduler consider several parameters for training such as age, academic year, past education performance, discipline, number of courses, and gap between two exams. The trained RL agent then able to provide examination schedule to a student depending on a student personal record, interests and abilities. The preliminary results are encouraging and more research would bring useful contribution of AI in various aspects of learning process of a student.

I. INTRODUCTION

Innovative educational approaches that involve the adapting of the educational process to the individual needs, interests and skills of each student can be defined as Personalized Learning [1]. The growing use of technologies like Artificial Intelligence (AI) and Machine Learning (ML), personalized learning allows to create customized learning experiences with the aim of increasing motivation, creating greater involvement, and improving the final result [2].

Widely recognized as the most disruptive innovation of recent years, AI is expanding its diffusion into an ever-increasing number of sectors: AI algorithms in healthcare are able to diagnose diseases through image analysis with high levels of accuracy and to create personalized treatment plans [3], [4]; in finance and smart offices, they allow to analyze large volumes of financial data to uncover trends and related investment opportunities [5]; self-driving cars and smart cities aspects are transforming mobility aiming to ensure better efficiency, sustainability and safety in the transport sector [6]; multi antenna communication including 5G and beyond networks [7], wireless sensor networks [8], dynamic treatment regimes [9], pervasive computing [10] and other significant progress many diverse areas including risk management in nuclear medicine department [11].

In the education sector, the concept of personalization goes beyond the simple adaptation of teaching materials to the needs of individual students, it includes an approach that recognizes and accommodates each student's unique learning style, interests and strengths. By personalizing the learning experience, teachers are able to create a more engaging and effective learning environment [12].

One aspect of personalization in education is customizing teaching to students preferred learning modes: some students may excel at visual learning, while others prefer auditory learning or other ways; by broadening the scope to additional teaching methods and resources, educators can ensure that all students have the opportunity to learn in the ways that best suit them. This personalized approach not only improves students' understanding, but also promotes a sense of autonomy in their educational journey [13].

Moreover, personalization in education can address students' individual academic needs and goals: for example, students can have different levels of proficiency in different subjects and thanks to personalized learning plans it is possible to identify specific areas where additional support is required. By thus aligning instruction with students' learning goals, educators make it easier for students to achieve academic goals at their own pace and avoid failure [14].

Another fundamental element, in addition to personalization, in maximizing the student learning process is represented by planning. Effective planning involves optimizing the allocation of time and resources to ensure that students have sufficient opportunities to learn and grow. This includes balancing teaching time with independent study, allocating time for collaborative activities and projects, and integrating breaks and reflection periods into the learning schedule.

In this context, the use of reinforcement learning (RL) algorithms in academic data favors innovative approaches to: personalizing and optimizing the educational experience for students, adapting materials and content based on individual student progress and preferences; optimization of allocation of educational resources, such as time, personnel, and materials. adaptive assessment systems

that can dynamically adjust the difficulty and format of assessments based on student performance;

intelligent tutoring systems, developing systems that provide personalized guidance and feedback to students and dynamically adjust strategies and activities.

In summary, the use of AI algorithms favors innovative approaches to personalize and optimize the students' educational experience; our contribution concerns the context of tutor systems to support students and in particular the implementation of a Reinforcement Learning based Intelligent System for personalized exam schedule with the aim of reducing exam fail.

II. TECHNICAL BACKGROUND

The AI, Artificial Intelligence, can be defined as the ability to develop intelligence on programmable machines with the aim of imitating the human brain [15]. Machine Learning (ML) is a subfield of AI, which concerns the question of how to implement software agents that automatically improve with experience.

Machine Learning is divided into three categories such as supervised learning, unsupervised learning and reinforcement learning [16].

Supervised learning is a type of learning based on the training dataset, a set of labeled input data, data for which the correct output is known. Training dataset is provided by a domain expert, who takes on the role of external supervisor of the process. The main goal of supervised learning is to build models enable of generalizing the relationship between inputs and the corresponding outputs in order to make the most accurate predictions on data not yet seen, based on its previous training.

The second category, Unsupervised Learning, is based on an appropriate study of the dataset with the aim of extracting knowledge and hidden patterns, there is no supervisor.

In Reinforcement Learning (RL), the entire learning process is linked to a specific objective. An agent interacts with an unknown environment, according to the try and error scheme; just the same way children learn it makes actions and observes what it happens [17]. Following the actions taken, the agent will receive feedback from the environment respectively in terms of reward for positive actions and penalties for negative actions. Thanks to this feedback it trains and acquires experience and knowledge about the environment.

RL involves problems of finding the optimal action to take in various scenarios to maximize the cumulative reward. The RL agent must develop a strategy (a comprehensive correlation between scenarios and actions) by experimenting with actions independently, without guidance from domain experts, like many other machine learning approaches.

Another crucial aspect of RL problems is the constant trade-off between exploiting the agent's existing knowledge of the environment (repeating actions previously taken in a given scenario) and exploring new actions that haven't been attempted in that scenario before.

A first major distinction that can be made between the RL

algorithms to use is that between single-agent and multi-agent, each with distinct applications and benefits. Single-Agent RL refers to scenarios in which a single agent interacts with the educational environment to optimize an individual student's learning path. In this context, the agent continuously analyzes the student's responses and performance to adapt the teaching content and improve the effectiveness of teaching, for example an intelligent tutor using Q-Learning or Deep Q-Networks (DQN) to personalize questions based on the student's answers. The goal is to maximize student progress by dynamically adjusting the difficulty of problems so that the student remains challenged but not overwhelmed. The benefit of this approach is the ability to create highly personalized learning paths, optimizing interaction with each student to improve their understanding and retention of the material.

Multi-Agent RL, on the other hand, involves multiple agents interacting with each other within the educational environment [18]. These agents can represent different students, tutors, or even different components of a complex educational system. Agents can collaborate or compete to achieve common or individual educational goals, for example an educational platform that uses multiple agents to simulate a collaborative learning environment, where students work together to solve complex problems.

Algorithms such as Multi-Agent Deep Deterministic Policy Gradient (MADDPG) can be used to coordinate the actions of agents so that each contributes effectively to the collective task. This approach facilitates collaborative and competitive learning, allowing students to benefit from interaction with their peers. Furthermore, it can improve virtual classroom dynamics by providing personalized support in a group learning context.

While single-agent RL focuses on optimizing the individual learning path, multi-agent RL supports both personalization by fostering group dynamics and promoting collaborative learning. However, multi-agent systems tend to be more complex to implement and manage than single-agent systems, as they require coordination and management of interactions between multiple agents. Finally, single-agent RL is ideal for personalized tutors and one-on-one learning assistants, while multi-agent RL is better suited for interactive and collaborative learning environments, such as classroom simulations and game-based learning platforms.

Q-Learning: This algorithm can be used to develop educational tutors who learn which actions (such as posing a question or reviewing a topic) lead to the best learning outcomes for students. Through a process of trial and error, the tutor updates a Q-table, which represents the value of each action in a given state, allowing the system to select the optimal action at any time.

Deep Q-Networks (DQN): DQN combines Q-Learning with deep neural networks to handle complex, continuous state space environments, such as those found in advanced e-learning platforms. This allows you to analyze a large amount of student data to make precise predictions on which educational content to propose next, optimizing learning in an

adaptive and personalized way.

Proximal Policy Optimization (PPO): This RL algorithm is particularly useful for managing stochastic policies in dynamic environments. In education, PPO can be used to create systems that not only decide which exercises or materials to propose, but also how to vary the difficulty and type of feedback based on the user's reaction in real time. PPO allows you to manage variations in teaching strategy in a more stable and efficient way, quickly adapting to the changing needs of students

III. RELATED WORKS

In this section, we will analyze the scientific literature regarding the use of AI and the development of adaptive systems in e-learning environments and the reinforcement learning approaches used to identify and model learning in education.

The authors of [1] [14] proposed the use of applications of the RL method in schools to improve learning activities and provide comfortable learning. The first system is personalized and adaptive e-learning platform based on Deep Q- Network RL and an online rule-based decision making implementation for elementary school. The second system is based on an intelligent educational environment for higher education in which the algorithm analyzes the change in student behavior, adapting teaching materials to improve the overall learning efficiency. In both, the intelligent learning created combines with the advantages of e-learning such as interactions, flexibility and experience.

In another work [12], authors presented a comparison between two models for customizing sequences of learning resources in a massive online open course (MOOC). The first model suggests sequences of learning resources that have been successful in the past, using case based reasoning and the Euclidean distance; while the second model recommends optimal learning resource sequences using Reinforcement's Q-Learning algorithm. A key role in the design of teaching resources is represented by the student's level of knowledge with respect to the level of complexity of the teaching resource with the aim of optimizing the MOOC learning process.

In the work of [13], the authors designed an intelligent adaptive e-learning system, based on RL that places particular importance on the adequacy of the student's real profile and its update compared to the one used in the learning path recommendation. The objectives of this system are the creation of a real profile of a given student, through the implementation of K-means and linear regression, and the recommendation of adaptive learning paths according to this profile, implementing the Q-learning algorithm.

Similarly, intelligent ambient assisted living systems are developed in [19] and [20]. The personalized systems are based on AI-powered tools like RL but are basically designed to provide personalized medication assistance and recommendation to the patients with cognitive impairments instead of students.

The next two quotes concern the use of gamification in the




Fig. 1. System Model

learning sector, in particular, the authors of the article [21] have demonstrated that gamification can be used to encourage student activities, increase engagement and evaluate their success; while the study [22] presents SEP-CyLE (Software Engineering and Programming Cyberlearning Environment), an online gamified tool designed to provide additional IT content to students.

In the article [23], to promote student involvement, satisfaction and performance, an integrated approach is developed that combines artificial intelligence with the result of the analysis of learning feedback through experimental research conducted to study the effects of learning.

The contribution of the authors of the article [24] is the development of an adaptive e-learning system capable of generating learning paths adapted to the profile of the individual learner. An approach is then proposed to dynamically compose adaptive online learning courses based on student activities, learning objectives and instructional design strategies using Q-learning. The algorithm gains knowledge by analyzing student behavior and provides the course content needed to achieve learning objectives based on positive and/or negative student feedback.

The work in [25] delves into the concept of an AI tutor that provides personalized learning paths and round-the-clock support to students. The AI tutor uses sophisticated algorithms and ML techniques to analyze a student's strengths, weaknesses, and learning style. By collecting data from various sources, such as assessments, quizzes and user interactions, the AI tutor ensures that students receive content and exercises proportionate to their individual progress promoting greater learning effectiveness.

IV. SYSTEM MODEL

This section presents the proposed work by introducing each component of the system and their corresponding functionality as shown in Figure 1.

The first step is to retrieve data such that the system accesses the university database to retrieve the necessary information for each student, which includes: personal data like age, working status, qualifications, skills and preferences, and academic data like courses attended, exams taken, grades obtained, credits acquired.

We model this scenario as Markov Decision Process (MDP)

problem and use RL to solve this MDP. We describe the components of modelled MDP as:

the **Environment** in the proposed framework is the entire university ecosystem, students and provided courses;

the **Agent** is our personalized exam scheduler tutor such that the agent has to learn the optimal exam schedule for a student;

the **State** represents the current status of the student with all information retrieved such that all the possible combinations and conditions are considered as states and RL agent has to choose an action (exam schedule) for a student;

the **Actions** are the possible actions that the system recommend to the students such that all the combinations that are available to a student to make selection for his/her examination;

the **Reward** is determined by the success or failure of the recommended actions, for example, passing the exam with a good grade is a reward, exam failure but also not showing up to the exam is a penalty for the system;

the **Policy** is the optimal strategy that a RL agent has to learn for a given environment. In our case, the RL agent (exam scheduler) has to learn the optimal exam schedule for each students according to each student existing credentials.

Our implementation produces a continuous cycle of learning and updating, based on student feedback, starting from the observation of the current status, in which the system collects the most recent information on the student's academic progress, the system analyze it with the current policy and recommends the next action (exam to book).

After the student has followed the recommendation, the system get feedback on the outcome (positive or negative outcome) and then make updates: the student's status based on the feedback received, and the policy update; the reinforcement learning model updates the policy based on feedback to improve future recommendations.

Moreover, we also consider privacy and security aspects and ensure that all student personal and academic data is treated with the utmost confidentiality and compliance with privacy regulations; and personalization. We consider that each student has different needs, so the system must be flexible to adapt to the peculiarities of each student.

V. DISCUSSION

In this section, we discuss the effectiveness of our reinforcement learning framework. We started our experimental activity with a data subset comprising all the target university courses: two years Master's Degree in Applied Behavioral and Cognitive Psychology, Business Administration, Organization and Management of Services for Sport and Physical Activities; three years Bachelor's Degree in Transport Science and Technology, Psychological Science and Techniques, Educational Science, Business Law and Economics, Computer Engineering; five years Master's Degree in Law.

The objective of the experimental phase is to demonstrate the extent to which the intelligent tutor improves each student's

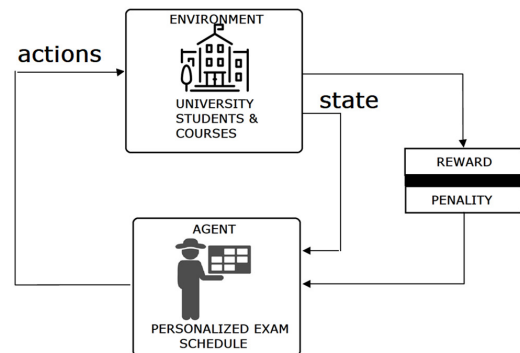


Fig. 2. Reinforcement Learning problem Model

learning path, facilitating final success and reducing the risk of failure.

By implementing a suitable RL algorithm, we aim to establish an optimal learning policy that minimizes the total number of learning steps. The reward function of the algorithm is structured to increase as the total number of learning steps decreases, thereby incentivizing more efficient learning performance.

Our evaluation will compare the outcomes of students who utilized the recommendations produced by our intelligent tutor with those who did not and we will employ a specifically defined metric that considers the relationship between the time taken and the results obtained. This comparison will help us understand the impact of our intelligent tutor on students' academic performance, specifically in terms of their ability to achieve learning objectives more efficiently and with a higher success rate.

To further substantiate our findings, we plan to conduct a series of statistical analyses, including hypothesis testing and confidence interval estimation, to determine the significance of the observed differences. We will also perform a detailed breakdown of the learning steps and rewards accrued by students in both groups, providing deeper insights into how the intelligent tutor influences learning behaviors and outcomes.

In addition to quantitative metrics, we intend to collect qualitative feedback from students and instructors to gauge their satisfaction and perceived effectiveness of the intelligent tutor. This approach will ensure a comprehensive evaluation of our reinforcement learning framework, highlighting both its strengths and areas for improvement.

Ultimately, the goal is to demonstrate that our intelligent tutor can serve as a valuable tool in educational settings, enhancing students' learning experiences and outcomes by providing personalized and effective learning recommendations.

As a case study in figure 3, for example, a student of age 25, taking a single course, and the difference between course enrollment and exam date is 164 days. The student followed a schedule where he attained 23 marks out of 30 and the

Age	AGE_NORM	Grade	GRADE_NORM	Ac_year	Univ_start_date	Exam_date	Date_diff	DIFF_NORM
40	0,9305	27	0,4669	2022/2023	24/10/2022	03/04/2023	161	0,072
22	0,2666	27	0,4669	2022/2023	10/10/2022	03/07/2023	266	0,1193
33	0,9454	28	0,607	2022/2023	29/08/2022	03/07/2023	308	0,1382
45	0,8377	30	1	2022/2023	14/11/2022	03/07/2023	231	0,1036
25	0,4901	23	0,184	2018/2019	22/10/2018	04/04/2019	164	0,0734
40	0,9305	27	0,4669	2021/2022	09/11/2021	04/04/2022	146	0,0653
53	0,4586	27	0,4669	2021/2022	01/03/2022	04/04/2022	34	0,0149
54	0,4139	24	0,3399	2021/2022	02/08/2021	04/04/2022	245	0,1099

Fig. 3. Data

proposed system should produced a schedule based on which student can achieve higher marks than 23 marks in the said experiment.

VI. CONCLUSION

Personalized learning has been demonstrated to be a useful idea in the education of a student. AI which has changed many facets of our lives has also been highly used in academia. RL is one of the intriguing AI approaches considered as an excellent tool for designing personalized solutions in education. RL algorithms can effectively take into account the student's characteristics. In this article, we explored the potential of utilizing RL for personalized exam scheduling. Our proposed framework considered many parameters including students' age, domain, number of courses, academic history, and the intermission between two consecutive exams. The trained RL agent will efficiently provide exam schedules based on their respective data. The initial results are inspiring and more research on this field would attain more valuable contributions in various aspects of personalized learning systems.

ACKNOWLEDGMENT

This research has been partially funded by the Regione Campania with the "PSR Campania 2014-2022 programme, Misura 16, Tipologia d'intervento 16.1.2." Project: "EVOOLIO - L'Evoluzione dell'Olio EVO Sannita tracciato con la Blockchain.

REFERENCES

- [1] W. Sayed, A. Noeman, A. Abdellatif, M. Abdelrazek, M. Badawy, A. Hamed, and S. El-Tantawy, "Ai-based adaptive personalized content presentation and exercises navigation for an effective and engaging e-learning platform," *Multimedia Tools and Applications*, vol. 82, pp. 1–31, 06 2022.
- [2] M. Naeem, S. T. H. Rizvi, and A. Coronato, "A gentle introduction to reinforcement learning and its application in different fields," *IEEE access*, vol. 8, pp. 209 320–209 344, 2020.
- [3] M. Naeem and A. Coronato, "An ai-empowered home-infrastructure to minimize medication errors," *Journal of Sensor and Actuator Networks*, vol. 11, no. 1, p. 13, 2022.
- [4] M. Naeem, A. Coronato, and G. Paragliola, "Adaptive treatment assisting system for patients using machine learning," in *2019 sixth international conference on social networks analysis, management and security (SNAMS)*. IEEE, 2019, pp. 460–465.
- [5] A. Coronato, G. De Pietro, and M. Esposito, "A semantic context service for smart offices," in *2006 International Conference on Hybrid Information Technology*, vol. 2. IEEE, 2006, pp. 391–399.
- [6] Z. Ullah, M. Naeem, A. Coronato, P. Ribino, and G. De Pietro, "Blockchain applications in sustainable smart cities," *Sustainable Cities and Society*, p. 104697, 2023.

- [7] M. Naeem, S. Bashir, Z. Ullah, and A. A. Syed, "A near optimal scheduling algorithm for efficient radio resource management in multi-user mimo systems," *Wireless Personal Communications*, vol. 106, no. 3, pp. 1411–1427, 2019.
- [8] A. Testa, A. Coronato, M. Cinque, and J. C. Augusto, "Static verification of wireless sensor networks with formal methods," in *2012 Eighth International Conference on Signal Image Technology and Internet Based Systems*. IEEE, 2012, pp. 587–594.
- [9] S. I. H. Shah, A. Coronato, M. Naeem, and G. De Pietro, "Learning and assessing optimal dynamic treatment regimes through cooperative imitation learning," *IEEE Access*, vol. 10, pp. 78 148–78 158, 2022.
- [10] M. Bakhouya, R. Campbell, A. Coronato, G. d. Pietro, and A. Ranganathan, "Introduction to special section on formal methods in pervasive computing," pp. 1–9, 2012.
- [11] S. I. H. Shah, M. Naeem, G. Paragliola, A. Coronato, and M. Pechenizkiy, "An ai-empowered infrastructure for risk prevention during medical examination," *Expert Systems with Applications*, vol. 225, p. 120048, 2023.
- [12] A. Flores, L. Alfaro, J. Herrera Quispe, and E. Cardenas, "Proposal models for personalization of e-learning based on flow theory and artificial intelligence," *International Journal of Advanced Computer Science and Applications*, vol. 10, 01 2019.
- [13] R. Mustapha, G. Soukaina, Q. Mohammed, and A. Es-Sâadia, "Towards an adaptive e-learning system based on deep learner profile, machine learning approach, and reinforcement learning," *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 5, 2023. [Online]. Available: <http://dx.doi.org/10.14569/IJACSA.2023.0140528>
- [14] S. Fu, "A reinforcement learning-based smart educational environment for higher education," *International Journal of e-Collaboration*, vol. 19, pp. 1–17, 01 2023.
- [15] U. b. Khalid, M. Naeem, F. Stasolla, M. H. Syed, M. Abbas, and A. Coronato, "Impact of ai-powered solutions in rehabilitation process: Recent improvements and future trends," *International Journal of General Medicine*, pp. 943–969, 2024.
- [16] M. Jamal, Z. Ullah, M. Naeem, M. Abbas, and A. Coronato, "A hybrid multi-agent reinforcement learning approach for spectrum sharing in vehicular networks," *Future Internet*, vol. 16, no. 5, p. 152, 2024.
- [17] M. Fiorino, M. Naeem, M. Ciampi, and A. Coronato, "Defining a metric-driven approach for learning hazardous situations," *Technologies*, vol. 12, no. 7, p. 103, 2024.
- [18] M. Naeem, A. Coronato, Z. Ullah, S. Bashir, and G. Paragliola, "Optimal user scheduling in multi antenna system using multi agent reinforcement learning," *Sensors*, vol. 22, no. 21, p. 8278, 2022.
- [19] A. Coronato and G. Paragliola, "A structured approach for the designing of safe aal applications," *Expert Systems with Applications*, vol. 85, pp. 1–13, 2017.
- [20] M. Ciampi, A. Coronato, M. Naeem, and S. Silvestri, "An intelligent environment for preventing medication errors in home treatment," *Expert Systems with Applications*, vol. 193, p. 116434, 2022.
- [21] F. Portela, "Towards an engaging and gamified online learning environment—a real casestudy," *Information*, vol. 13, no. 2, 2022.
- [22] L. Zahedi, J. Batten, M. Ross, G. Potvin, S. Damas, P. Clarke, and D. Davis, "Gamification in education: a mixed-methods study of gender on computer science students' academic performance and identity development," *Journal of Computing in Higher Education*, vol. 33, pp. 1–34, 08 2021.
- [23] F. Ouyang, M. Wu, L. Zheng, L. Zhang, and P. Jiao, "Integration of artificial intelligence performance prediction and learning analytics to improve student learning in online engineering course: Revista de universidad y sociedad del conocimiento," *International Journal of Educational Technology in Higher Education*, vol. 20, no. 1, p. 4, Dec 2023 2023.
- [24] C. W. Fernandes, T. Miari, S. Rafatirad, and H. Sayadi, "Unleashing the potential of reinforcement learning for enhanced personalized education," in *2023 IEEE Frontiers in Education Conference (FIE)*, 2023, pp. 1–5.
- [25] A. Robert, "Your ai tutor: Personalized learning paths and 24/7 support," 04 2024.

Disease Diagnosis On Ships Using Hierarchical Reinforcement Learning

Farwa Batool
Quaid-i-Azam University
Islamabad Pakistan
Email: farwabatool@ele.qau.edu.pk

Tehreem Hasan
Quaid-i-Azam University
Islamabad Pakistan
Email: tehreemhasan@ele.qau.edu.pk

Giancarlo Tretola
Department of Computer Engineering
Università Giustino Fortunato
Benevento Italy
Email: g.tretola@unifortunato.eu

Zaib Ullah
Department of Computer Engineering
Università Giustino Fortunato
Benevento Italy
Email: z.ullah@unifortunato.eu

Musarat Abbas
Quaid-i-Azam University
Islamabad Pakistan
Email: mabbas@qau.edu.pk

Abstract—Every year about 30 million people travel by ship worldwide often in extreme weather conditions and polluted environments and many other factors that impact the health of passengers and crew staff. Such issues require medical staff for passenger health care. We introduce a model based on Reinforcement learning (RL) which is used in the dialogue system. We incorporate the Hierarchical reinforcement learning (HRL) model with the layers of Deep Q-Network for dialogue oriented diagnosis system. Policy learning is integrated as policy gradients are already defined. We created a two-stage hierarchical strategy. We used the hierarchical structure with double-layer policies for automatic disease diagnosis. A double layer means it splits the task into sub-tasks named high-state strategy and low-level strategy. It has a user simulator component that communicates with the patient for symptom collection low-level agents inquire about symptoms. Once it's done collecting it sends results to the high-level agent which activates the D-classifier for the last diagnosis. When it's done its sent back by the user simulator to patients to verify the diagnosis made. Every single diagnosis made has its reward that trains the system

I. INTRODUCTION

MARITIME TRANSPORTATION plays a vital role in global trade and passenger transport contributing to economic development and connectivity [14]. Maritime transportation is the backbone of global trade, as ships carry over 80 percent of trading goods worldwide [34]. Almost every industry is changing due to technology and new methods of operation, but the maritime sector is currently seeing this transition most quickly [26]. Further investigation provides insights into the function of innovative communications technology, including virtual telemedicine and secure radio expertise, and assesses their practicality in the context of emergency maritime medicine [8], [12]. There is always a need of medical facilities for passengers and crew members. One of the biggest challenge in it is timely and accurate diagnosis of disease. As ships have limited resources and lack of medical staff on board so we can not rely on traditional methods. So we move towards Machine learning and Artificial Intelligence

(AI) to train system to do automatic diagnosis [2]. AI has emerged as a revolutionary force in many field like 5G vehicular networks [10], rehabilitation [24], MIMO communication [17] and also in healthcare, offering new methods to the way we do disease identification, its treatment, and tracking. The implementation of AI in healthcare is enhancing diagnostic accuracy [15]. Specially ,Hierarchical reinforcement learning (HRL) is a promising method to extend traditional reinforcement learning to solve more complex tasks [38]. Hierarchical reinforcement learning (HRL) provides more broad spectrum to RL, by offering a divide-and-conquer methodology. In this methodology, the intricate and challenging problems, are divided into multiple smaller problems. These divided problems are easier to solve and their solutions can be regenerative to solve other related problems. This methodology has preceding been successfully used to speed up many offline preparing and organising algorithms where the variables of the environment are known in advance [7]. Hierarchical reinforcement learning (HRL) is a layered algorithm based on RL. HRL has been evidenced to be efficient in challenges with deferred and infrequent rewards and minimizing the learning difficulty by splitting the long-term goal into stages [35]. The symptom collection process of multiple phases of consultation between the agent and the patient as a Markov decision process, and uses the reinforcement learning algorithm for training [30]. our contribution is implementing the HRL by assigning rewards to correct symptom query in result of agent collecting the symptom and relating it with certain disease. policy learning is integrated as policy gradients are already defined. As we are using hierarchical reinforcement learning it creates two stage hierarchical strategy, fist stage is high level strategy which triggers the low level strategy. Low level strategy have multiple agents working as symptoms checkers and disease classifiers. Each Agent is responsible for investigating certain types of diseases. At the end we have disease classifier which is responsible to check responses from all agents and conclude

disease diagnosed. Every disease have relation with symptoms and symptoms are also related with more than one disease. So for achieving maximum accuracy its necessary to understand symptoms and narrow down options of diseases at every single question with dialogue simulator. Now on ships as we have limited medical staff so its doing diagnosis using HRL, in which we have Agents every single agent is specialized for certain field providing broad spectrum of diseases to be diagnosed. The paper organised as follows, firstly we have related work. As Reinforcement learning specifically hierarchical reinforcement learning is emerging and is popular for classification, So we mentioned worked done earlier. Secondly proposed framework model is which explains all components in the model that includes leader, agent, user simulator, d-classifier. Its shown in detail in figure 1. Thirdly we have benchmark models which describe all the best models we are comparing with. Lastly we have results and conclusions.

II. RELATED WORK

This section outlines some related works on the use of reinforcement learning for healthcare problems.

Dynamic Treatment Regime (DTR) is has an importance in healthcare as well as for medical research. DTR are considered as sequence of alternative treatment paths and any of these treatments can be adapted depending on the patient's conditions [6]. Therefore, the authors in [22] apply a cooperative imitation learning approach to utilize information from both negative and positive trajectories to learn the optimal DTR. The given framework minimizes the chance of choosing any treatment that results in a negative outcome during the medical examination. However, the proposed work is not suitable to employ for the disease diagnosis on ships.

Online symptom checkers by [20] have been put into action to recognise the possible causes and treatments for diseases based on a patient's symptoms. The work in [11] uses deep RL for fast disease diagnosis. Similarly, authors in [25] utilize an approach of automatic development of a dialogue manager capable of doing goal-oriented dialogues for the health domain. While the work in [29] employs a hierarchical RL is used for automatic captioning the video.

A machine learning method upper confidence bound is utilized in [16] to assist patients during their medication process at home. Authors considered the cognitive and physical impairments of the patients in the training of the machine learning model. A similar work is also done in [5] but with the help of Thompson sampling method. However, these systems are useful to specific scenarios during medication at home.

An end to end multi-channel conversational interface for dynamic and co-operative target setting is developed in [29], which integrates collective reward (task/persona/sentiment) for task success, personalized augmentation and user-adaptive behavior. Furthermore, an automatic diagnostic system is designed in [27] by applying both evident and inherent symptoms utilized by the Deep-Q Network Reinforcement Policy.

Moreover, there are some AI based solutions for the continuous and remote monitoring of unpredictable health issues.

Such a failure mode and effect analysis is given in [4] and [3] for a specific mobile health monitoring system. Both of these systems were designed to provide remote healthcare solutions but these are for certain cases and environments and cannot be generalised for other cases.

The works in [19] and [23] use AI techniques for risk management in nuclear medication department. The later will is the extension of former one and discuss the risk cases during examination at such departments. Although, the proposed systems are useful to avoid possible risk at nuclear medication departments but are not useful for healthcare solutions at ships. an End-to-End Knowledge-routed Relational Dialogue System (KR-DS) that enables dialogue management, natural language understanding, and natural language generation to cooperatively optimize via reinforcement learning is presented in [1]. [32]. Q-learning algorithm is used in [18] to create an optimal controller for cancer chemotherapy drug dosing. Major depressive disorder treatment is considered in [21]. The authors have utilized the strong transfer ability of HRL to build a cross-domain dialogue system, which learned shareable information in similar subdomains of different main domains to train a general underlying policy.

Hybrid and hierarchical RL methods gained significant attention in recent years [10]. The proposed work presents extended RL structure as hierarchical structure that has two-stage policies for automatic diagnosis. it has hierarchical structure with double layer policies for automatic disease diagnosis. Double layer means it splits the task into sub-tasks named as high-state strategy and low level strategy. User simulator communicates with patient for symptom collection low level agent inquire symptoms. Once its done collecting it sends results to high level agent which activates the D-classifier for last diagnosis. When its done its send back by user simulator to patients to verify diagnosis made.

III. MODEL FRAMEWORK

The disease diagnosis model finds the policy π for the maximum reward. For disease diagnosis Markov decision process is used in which $M = [S, A, P, \gamma]$ [9]. S is the state, S^h is state in high stage strategy, S^{li} is state in low-stage strategy, n is the number of low strategy agents. All states can be expressed as $S = S^h \cup \{S^{li}\}_{i=1}^n$. For actions, A^h is high stage agent's action, A^{li} is low stage action. n is the number of low strategy agents. All actions are expressed as $A = A^h \cup \{A^{li}\}_{i=1}^n$. All dialogue rewards is shown by R . State transition model is shown by P . γ is the discount rate used to compute Q value function. The major aim is to optimize Markov decision process $M = [S, A, P, \Gamma]$ and identify the policy π that elevate the cumulative discount reward for all (S, A) .

In this paper, we extend simple RL structure into hierarchical structure that has two-stage policies for automatic diagnosis. Framework is shown in the Figure 1 it is hierarchical structure with double layer policies for automatic disease diagnosis. Double layer means it splits the task into sub-tasks named as high-state strategy and low level strategy. Idea is inspired by hospital consultation in real world. It works in a way that

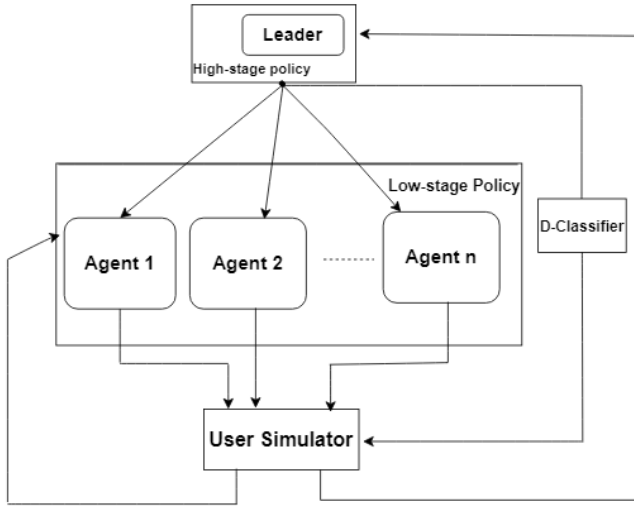


Fig. 1. Model Diagram

High-state Agent gets the current initial state as S_t , then it appoints a low level agent to communicate with user simulator for symptom collection. In Figure 1, it has four main parts Leader, Agent, User simulator and Disease classifier. Current initial state as S_t is encoded as a vector that depicts the level of each symptom and also about number of iterations necessary.

Consider a doctor that asks symptoms from patient. They will first consider that patient have certain disease and start asking related symptoms. Similar to that agent chooses a symptom to inquire the patient $A_t \in S$ The possible user responses could be (true/false/unknown). If a_t is element in set of diseases, agent will inform user about diagnosis made, and diagnosis is made dialogue session would end and accuracy depends on correctness of diagnosis.

A. Strategy of Leader Model

In leader model its main task is to figure out if its activating the D-classifier or the agent to collect more symptoms. Once the leader activates the agent it will interact with user N number of cycles (dialogue rounds) until sub task is terminated. For action a_t^l the reward of the leader is r_t^l . Γ is the discount factor and $r_{t+t'}^e$ is the reward given by user simulator to low stage agent for current cycle. One reward is generated for disease classifier shown as r_t^e . In formula d is the action to activate the agent A^i .

$$r_t = \begin{cases} \sum_{t'=1}^N \Gamma^{t'} r_{t+t'}^e, & \text{if } a_t^l = A^i \\ r_t^e, & \text{if } a_t^l = d \end{cases} \quad (1)$$

The rewards obtained from user simulator will be aggregated as the reward of High-stage agent which is the high-stage reward calculated in equation(1).

B. Strategy of Agent Model

The objective of agent is to optimize the expected cumulative discounted reward. For that we use bellman equation in it

Q-value function illustrates cumulative reward. In equation θ_l is the parameter of present policy network. Action of agent is shown by a_t^l and after taking action the next dialogue state is S_{t+1} to the policy π .

$$Q_t^\pi(s_t, a_t^l | \theta^l) = r_t^l + \mathbb{E}_{(S_{t+1}, a_{t+1}^l)} [\Gamma_t^T Q_t^\pi(S_{t+1}, a_{t+1}^l | \theta^l)] \quad (2)$$

The low-stage agent has task of compiling symptoms by taking to user simulator, which is activated by high-stage agent. The high level agent has layers of DQN and parameters of the network is shown by θ_l . The parameters keep updating in training by decreasing the mean-square error(MSE) between the Q-values of target network achieved and the Q-value of current one. That MSE is utilized as loss function of the advance policy network as shown in equation (3).

$$L(\theta^l) = \mathbb{E} [r_t^l + \Gamma_t^T \max_{a_{t+1}} Q_t^*(S_{t+1}, a_{t+1}^l | \theta^l) - Q_t^\pi(S_{t+1}, a_{t+1}^l | \theta^l)]^2 \quad (3)$$

In equation (3) first term is Q value of target network achieved and second one is Q value of present network.

C. User Simulator

The user simulator is the part of system that is responsible of communicating with agent and also contains the user aims in the data set. AT the start of every dialogue session it samples the aims randomly from training set. User aim hold two types of symptoms named as explicit and implicit symptoms. Explicit symptoms are provided to agent as initial input and with the help of that it will discover implicit symptoms while interacting with patient. During the interaction if it gets correct symptom then it will get reward as 1 , with incorrect symptom it will get reward of -1 and for an unknown symptom it will get reward of 0. Once its done collecting symptoms from patient low-stage agent activates high-stage agent and then the disease classifier for final classification of disease.

D. synthetic Dataset

On ships we have vast range of diseases that can occur, so having such big real world data set was almost impossible so we used synthetic data set available as Data/Fudan-Medical-Dialogue2.0 to show the effectiveness of HRL. In it every disease is linked with set of symptoms, not only that every single symptom has a probability for a certain disease. Now for identification process out of many symptoms in data set we choose any of explicit symptoms among those provided by the patient , that one symptom has more importance and rest of the symptoms are treated as implicit symptoms.

IV. BENCHMARK MODELS

First work is done on dialogue system which used task oriented disease diagnoses. It used one layer policy structure based DQN wich is called FLAT-DQN it has to do with choosing actions in each turn of dialogues[30]. After that there is a dialogue system for automatic medical diagnosis that communicates with patients to collect extra symptoms other than their self-reports and do automatic diagnose. It

uses KR-DS that treats all diseases and all symptoms equally [33]. HDNO, a hierarchical reinforcement learning model, to improve performance and is validated on dialogue-based MultiWoz datasets [28]. HRL is a hierarchical reinforcement learning model which uses disease classifier for classification of symptoms separately [13]. GAMP, a model that integrates the generative adversarial network(GAN). Its policy was also DQN based used generator to generate action and a discriminator is there to check if its a good action taken on base of reward achieved [31]. HRL-pre-T , Its Hrl pre trained has two levels of policy just like us but one visible difference is that it trains the models separately and we train them together [11]. HRL is the model that used both real world and synthetic dataset and used in disease diagnosis [36]. KN-HRL is the enhanced model that creates the disease symptom relation matrix and do disease diagnosis based on patient’s utterances [37].

V. RESULTS AND CONCLUSIONS

In order to check performance of our model we conduct experiment on same synthetic dataset. We did comparison of all models that includes Flat-DQN, KR-Ds, REFUEL, GAMP, HRL-Pre-T, KNHRL and Lastly our HRL model. Flat-DQN, KR-Ds, REFUEL performed almost similarly. Flat-DQN, KR-Ds are good models but performed best with the short dialogues. KNHRL and Lastly HRL model performed well but with less accuracy with larger dialogues. We present a comparison of all as given in Table 1. HRL(ours) used publicly available data set with the more medical knowledge in format of dialogues. It used disease symptom relation and symptom disease for training and testing both , also multiple rounds of dialogues with user simulator and patient and multiple layers of DQN which improves accuracy.

Table 1

	Test Accuracy	Avg turns	Match rate
Flat-DQN	0.343	1.23	0.023
KR-Ds	0.357	6.24	0.388
REFUEL	0.416	4.56	0.161
GAMP	0.409	1.36	0.077
HRL-Pre-T	0.452	6.838	/
HRL	0.504	6.48	0.495
KNHRL	0.558	20.98	0.333
HRL (ours)	0.627	3.00	0.506

In future work we hope to gain more accuracy and collect some real world dataset. We think that with further more improvements this model can solve the problem of shortage of medical staff in the entire world.

ACKNOWLEDGMENT

This project has been partially funded by the “Programma Nazionale Ricerca, Innovazione e Competitività per la transizione verde e digitale 2021/2027 destinate all’intervento del FCS “Scoperta imprenditoriale” - Azione 1.1.4 “Ricerca collaborativa” - with the project SIAMO (Servizi Innovativi per

l’Assistenza Medica a bOrdo) project number F/360124/01-02/X75.

REFERENCES

- [1] Qanita Bani Baker, Safa Swedat, and Kefah Aleesa. Automatic disease diagnosis system using deep q-network reinforcement learning. In *2023 14th International Conference on Information and Communication Systems (ICICS)*, pages 1–6, 2023.
- [2] Mohamed Bakhouya, Roy Campbell, Antonio Coronato, Giuseppe de Pietro, and Anand Ranganathan. Introduction to special section on formal methods in pervasive computing, 2012.
- [3] Marcello Cinque, Antonio Coronato, and Alessandro Testa. Dependable services for mobile health monitoring systems. *International Journal of Ambient Computing and Intelligence (IJACI)*, 4(1):1–15, 2012.
- [4] Marcello Cinque, Antonio Coronato, and Alessandro Testa. A failure modes and effects analysis of mobile health monitoring systems. In *Innovations and advances in computer, information, systems sciences, and engineering*, pages 569–582. Springer, 2012.
- [5] Antonio Coronato and Muddasar Naeem. A reinforcement learning based intelligent system for the healthcare treatment assistance of patients with disabilities. In *International Symposium on Pervasive Systems, Algorithms and Networks*, pages 15–28. Springer, 2019.
- [6] Antonio Coronato, Muddasar Naeem, Giuseppe De Pietro, and Giovanni Paragliola. Reinforcement learning for intelligent healthcare applications: A survey. *Artificial Intelligence in Medicine*, 109:101964, 2020.
- [7] Antonio Coronato and Giovanni Paragliola. A structured approach for the designing of safe aal applications. *Expert Systems with Applications*, 85:1–13, 2017.
- [8] Jonathan S Dillard, William Maynard, and Rahul Kashyap. The epidemiology of maritime patients requiring medical evacuation: a literature review. *Cureus*, 15(11), 2023.
- [9] Mario Fiorino, Muddasar Naeem, Mario Ciampi, and Antonio Coronato. Defining a metric-driven approach for learning hazardous situations. *Technologies*, 12(7):103, 2024.
- [10] Mansoor Jamal, Zaib Ullah, Muddasar Naeem, Musarat Abbas, and Antonio Coronato. A hybrid multi-agent reinforcement learning approach for spectrum sharing in vehicular networks. *Future Internet*, 16(5):152, 2024.
- [11] Hao-Cheng Kao, Kai-Fu Tang, and Edward Chang. Context-aware symptom checking for disease diagnosis using hierarchical reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [12] Umamah bint Khalid, Muddasar Naeem, Fabrizio Stasolla, Madiha Haider Syed, Musarat Abbas, and Antonio Coronato. Impact of ai-powered solutions in rehabilitation process: Recent improvements and future trends. *International Journal of General Medicine*, pages 943–969, 2024.
- [13] Kangerbei Liao, CHENG ZHONG, Wei Chen, Qianlong Liu, Baolin Peng, Xuanjing Huang, et al. Task-oriented dialogue system for automatic disease diagnosis via hierarchical reinforcement learning, 2021.
- [14] Luigia Mocerino, Fabio Murena, Franco Quaranta, and Domenico Toscano. Validation of the estimated ships’ emissions through an experimental campaign in port. *Ocean Engineering*, 288:115957, 2023.
- [15] Muddasar Naeem and Antonio Coronato. An ai-empowered home-infrastructure to minimize medication errors. *Journal of Sensor and Actuator Networks*, 11(1):13, 2022.
- [16] Muddasar Naeem, Antonio Coronato, and Giovanni Paragliola. Adaptive treatment assisting system for patients using machine learning. In *2019 sixth international conference on social networks analysis, management and security (SNAMS)*, pages 460–465. IEEE, 2019.
- [17] Muddasar Naeem, Antonio Coronato, Zaib Ullah, Sajid Bashir, and Giovanni Paragliola. Optimal user scheduling in multi antenna system using multi agent reinforcement learning. *Sensors*, 22(21):8278, 2022.
- [18] Regina Padmanabhan, Nader Meskin, and Wassim M. Haddad. Learning-based control of cancer chemotherapy treatment**this publication was made possible by the gsra grant no. gsra1-1-1128-13016 from the qatar national research fund (a member of qatar foundation). the findings achieved herein are solely the responsibility of the authors. *IFAC-PapersOnLine*, 50(1):15127–15132, 2017. 20th IFAC World Congress.
- [19] Giovanni Paragliola, Antonio Coronato, Muddasar Naeem, and Giuseppe De Pietro. A reinforcement learning-based approach for the risk management of e-health environments: A case study. In *2018 14th*

- international conference on signal-image technology & internet-based systems (SITIS)*, pages 711–716. IEEE, 2018.
- [20] Yu-Shao Peng, Kai-Fu Tang, Hsuan-Tien Lin, and Edward Chang. Refuel: Exploring sparse features in deep reinforcement learning for fast disease diagnosis. *Advances in neural information processing systems*, 31, 2018.
- [21] A.John Rush, Maurizio Fava, Stephen R Wisniewski, Philip W Lavori, Madhukar H Trivedi, Harold A Sackeim, Michael E Thase, Andrew A Nierenberg, Frederic M Quitkin, T.Michael Kashner, David J Kupfer, Jerrold F Rosenbaum, Jonathan Alpert, Jonathan W Stewart, Patrick J McGrath, Melanie M Biggs, Kathy Shores-Wilson, Barry D Lebowitz, Louise Ritz, George Niederehe, and for the STAR*D Investigators Group. Sequenced treatment alternatives to relieve depression (star*d): rationale and design. *Controlled Clinical Trials*, 25(1):119–142, 2004.
- [22] Syed Ihtesham Hussain Shah, Antonio Coronato, Muddasar Naeem, and Giuseppe De Pietro. Learning and assessing optimal dynamic treatment regimes through cooperative imitation learning. *IEEE Access*, 10:78148–78158, 2022.
- [23] Syed Ihtesham Hussain Shah, Muddasar Naeem, Giovanni Paragliola, Antonio Coronato, and Mykola Pechenizkiy. An ai-empowered infrastructure for risk prevention during medical examination. *Expert Systems with Applications*, 225:120048, 2023.
- [24] Beata Sokolowska, Wiktor Świdorski, Edyta Smolis-Bąk, Ewa Sokolowska, and Teresa Sadura-Siekłucka. A machine learning approach to evaluate the impact of virtual balance/cognitive training on fall risk in older women. *Frontiers in Computational Neuroscience*, 18:1390208, 2024.
- [25] Milene Santos Teixeira, Vinícius Maran, and Mauro Dragoni. The interplay of a conversational ontology and ai planning for health dialogue management. In *Proceedings of the 36th annual ACM symposium on applied computing*, pages 611–619, 2021.
- [26] Edvard Tijan, Marija Jović, Saša Aksentijević, and Andreja Pucihar. Digital transformation in the maritime transport sector. *Technological Forecasting and Social Change*, 170:120879, 2021.
- [27] Abhisek Tiwari, Tulika Saha, Sriparna Saha, Shubhashis Sengupta, Anutosh Maitra, Roshni Ramnani, and Pushpak Bhattacharyya. Multi-modal dialogue policy learning for dynamic and co-operative goal setting. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021.
- [28] Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. Modelling hierarchical structure between dialogue policy and natural language generator with option framework for task-oriented dialogue system. *arXiv preprint arXiv:2006.06814*, 2020.
- [29] Xin Wang, Wenhui Chen, Jiawei Wu, Yuan-Fang Wang, and William Yang Wang. Video captioning via hierarchical reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4213–4222, 2018.
- [30] Zhongyu Wei, Qianlong Liu, Baolin Peng, Huaixiao Tou, Ting Chen, Xuanjing Huang, Kam-fai Wong, and Xiangying Dai. Task-oriented dialogue system for automatic diagnosis. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–207, Melbourne, Australia, July 2018. Association for Computational Linguistics.
- [31] Yuan Xia, Jingbo Zhou, Zhenhui Shi, Chao Lu, and Haifeng Huang. Generative adversarial regularized mutual information policy gradient framework for automatic diagnosis. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):1062–1069, Apr. 2020.
- [32] Lin Xu, Lin Xu, Qixian Zhou, Qixian Zhou, , Ke Gong, Xiaodan Liang, Xiaodan Liang, Jianheng Tang, Jianheng Tang, Jianheng Tang, Lin Li, and Liang Lin. End-to-end knowledge-routed relational dialogue system for automatic diagnosis. *null*, 2019.
- [33] Lin Xu, Qixian Zhou, Ke Gong, Xiaodan Liang, Jianheng Tang, and Liang Lin. End-to-end knowledge-routed relational dialogue system for automatic diagnosis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7346–7353, 2019.
- [34] Ran Yan, Dong Yang, Tianyu Wang, Haoyu Mo, and Shuaian Wang. Improving ship energy efficiency: Models, methods, and applications. *Applied Energy*, 368:123132, 2024.
- [35] Qian Zhang, Tianhao Li, Dengfeng Li, and Wei Lu. A goal-oriented reinforcement learning for optimal drug dosage control. *Annals of Operations Research*, pages 1–21, 2024.
- [36] Cheng Zhong, Kangenbei Liao, Wei Chen, Qianlong Liu, Baolin Peng, Xuanjing Huang, Jiajie Peng, and Zhongyu Wei. Hierarchical reinforcement learning for automatic disease diagnosis. *Bioinformatics*, 38(16):3995–4001, 07 2022.
- [37] Ying Zhu, Yameng Li, Yuan Cui, Tianbao Zhang, Daling Wang, Yifei Zhang, and Shi Feng. A knowledge-enhanced hierarchical reinforcement learning-based dialogue system for automatic disease diagnosis. *Electronics*, 12(24), 2023.
- [38] Qijie Zou, Xiling Zhao, Bing Gao, Shuang Chen, Zhiguo Liu, and Zhejie Zhang. Relabeling and policy distillation of hierarchical reinforcement learning. *International Journal of Machine Learning and Cybernetics*, pages 1–17, 2024.

Dashboard User Interface (UI) Implementation for Remote Critical Infrastructure Inspection by using UAV/Satellite in times of Pandemic

Romaio Bratskas
0000-0002-4337-5828
Skyld Security and Defence Ltd,
Alkaiou 23, 2404, Nicosia,
Cyprus
Email: rb@skyld.com.cy

Dr. Dimitrios Papachristos
0000-0002-3453-0022
Skyld Security and Defence Ltd,
Alkaiou 23, 2404, Nicosia,
Cyprus
Email:
d.papachristos@skyld.com.cy

Dr. Petros Savvidis
0009-0000-0444-2314
Skyld Security and Defence Ltd,
Alkaiou 23, 2404, Nicosia,
Cyprus
Email:p.savvidis@skyld.com.cy

Dr. George Leventakis
Dept. Shipping, Trade &
Transport, Univ. of Aegean,
Korai St. 2A, 82132, Chios,
Greece
Email: glevantakis@aegean.gr

Enea Qerama
Skyld Security and Defence Ltd,
Alkaiou 23, 2404, Nicosia,
Cyprus
Email: e.qerama@skyld.com.cy

George Dahrouje
Skyld Security and Defence Ltd,
Alkaiou 23, 2404, Nicosia,
Cyprus
Email: d.dahrouje@skyld.com.cy

Abstract—In times of pandemic, many activities of the society, economy are minimized due to the risk of transmission. In particular, in the period of Covid-19, with the implementation of the Lockdown, many infrastructure monitoring and maintenance activities were suspended to prevent the spread of the virus. In essence, pandemics of highly contagious viruses may directly impact the critical infrastructure monitoring sector. More specific, in the context of monitoring critical infrastructure through satellites and UAVs, data processing involves extracting valuable insights, detecting potential threats, and assessing the overall condition of the infrastructure. This processed information is then used to make informed decisions regarding maintenance, security measures, and response strategies to mitigate risks and safeguard the critical assets. In this paper, we present a user interface dashboard dedicated to inspecting the critical infrastructure events captured from UAV or satellite. The design and architecture of the Dashboard User Interface its primary goal continues to be delivering real-time images to users, showcasing areas/components/points of failure in critical infrastructure, including damaged components, structural issues, corrosion, vegetation obstruction etc.

Index Terms—remote inspection, critical infrastructures, pandemic era, User Interface dashboard

I. INTRODUCTION

CURRENTLY, remote inspections are not only considered simply as a backup in case of major disruption (i.e. pandemic), but as a strategic and sustainable solution for the future. Critical Infrastructures (CIs) across industries have realized the benefits of harnessing new technology to streamline inspection processes, improve accessibility, and enhance overall operational efficiency. The SUNRISE platform aims

to provide users with real-time information on critical infrastructure by analyzing image and video feeds. This information includes details on damaged components, structural issues, corrosion, and obstructing vegetation. The tool utilizes AI-assisted components to process data from satellite and UAV feeds and present it in the user interface. In addition, the use of UAVs is important and more convenient than the use of satellites thanks to the technological evolution of UAVs (communications, propulsion/navigation, swarms etc.) providing high reliability in data recording [1][2][3][4].

In this context, the user interface (UI) tool of the SUNRISE platform for remote infrastructure inspection provides users with real-time images of areas, components, or points of failure in critical infrastructure, such as damaged components, structural issues, corrosion, obstructing vegetation, and more.

In this paper, we present a user interface dashboard tool dedicated to inspecting the critical infrastructure events captured from UAV or satellite and an architectural paradigm including back-end design. UI design is a complicated process that requires detailed analysis of operators' performance and preference. Furthermore, developments in technology require an understanding of "trust" aspects of interaction [5].

II. BACKGROUND

Critical infrastructure (CIs) is a critical component of the economy in all countries with technological development [6]. Critical infrastructure, as defined by [7], refers to systems that include industries, organisations with distribution & transportation capabilities that provide a continuous flow of essential services vital to the defence and economic security of society. Typically, these systems are also called lifeline systems.

¹This work was supported by SUNRISE H2020 (EU Project)

Critical infrastructure incorporates various systems such as electric power systems, telecommunications, water treatment and supply, natural gas supply, transportation systems, and healthcare systems [8]. The term critical infrastructure is defined in Section 1016(e) of the USA Patriot Act of 2001 as those “systems and goods, both physical and virtual, so vital to the nation that their malfunctioning or destruction would produce a debilitating impact on the security of citizens, on the economic security of the nation, on national public health and on any combination of the above” [9].

The CIs framework is constantly changing to reflect current concerns and respond to new issues, notably security and resilience [10]. At the same time, the quantity and variety of CIs have increased dramatically in recent decades [11]. Moreover, their protection from the various shocks and stresses, and ensuring their continuous functioning have become a basic priority [12]. In addition, reducing damages on CIs is critical in societal well-being and achieving resilience and sustainability [13].

One of the great challenges is to make CIs truly sustainable from all points of view: economic, social, environmental. CIs play a key role in directly and indirectly impacting progress on achieving sustainable development goals (SDGs). The recent pandemic crisis showed that the future is uncertain for humanity. Therefore, societies should prepare themselves to address the weaknesses that make us vulnerable to such important risks. The sustainability, reliability and security of CS are issues of huge importance with international resonance that require a new approach. Therefore, several techniques and methods (based in science & technology) have been developed to ensure CIs, taking into consideration different issues such as human error, maintenance policies, energy, water supply, healthcare protection and emergency transportation in case of disaster [14]. CIs are, therefore, those material resources, services, IT systems, networks and infrastructure assets that, if damaged or destroyed, would cause serious repercussions on the crucial functions of society, including the supply chain, health, security and the economic or social well-being of the state and the population [15].

That's why the use of technologies such as UXVs (i.e. UAV, UGV) are a new approach to inspecting critical infrastructure. For example, the use of UAVs has significant potential to increase efficiency, reduce O&M costs, improve operational safety as well as minimize production downtimes in the wind energy sector [16]. Essentially, UAVs are a useful data collection tool for capturing large volumes of data, such as high-resolution images and other information in a relatively short period of time. This information can be used to monitor progress or changes in a critical infrastructure for its sustainability and safety [4].

III. TOOL ARCHITECTURE

A. High Level Design

The Tool consists of the following parts (Fig.1): the Web Application-Dashboard UI (including a and b), the MQTT Bus and the Backend Coordinator (including a Data Base and

Reporting subsystem). As main functionalities, the user interface (UI) tools and dashboards furnish a dynamic web platform, empowering end-users to have a complete engagement with all inspection infrastructures' components of SUNRISE system. This module incorporates contemporary tools and presents a map where during inspection, constantly refreshed with real-time data on inspection point. This Tool enhances the inspection process, enabling end-users to accomplish a full range of inspect activities, receive event-driven messages, that means anomalies detection with use of AI algorithms-based methods.

More specific, the selected web application architecture is described through the numbered bus lines as follows:

1. Incoming messages/events from UAV/Satellite systems are received. All this data is routed through an MQTT bus system. Within this system, the data is systematically queued, ensuring a sequential flow.
2. The Backend Coordinator processes all incoming messages/events. It retrieves the data at the front of the MQTT queue.
3. All incoming messages/events are internally stored in the Backend Inventory (MongoDB server).
4. The Backend Coordinator sends live or historical data to the Dashboard UI for visualization and responds to historical data requests from the Dashboard UI.
5. The Dashboard UI communicates with the Google Maps infrastructure to render maps, markers, points of interest, and heat maps, among other elements.
6. The Backend Coordinator sends requests to the Reporting Subsystem in order to compile the requested data and then receives the results.
7. The Reporting Subsystem and the Backend Inventory communicate with each other in order to process the requests, and subsequently transmits the results to the Backend.
8. The Dashboard UI obtains an Access Token from the Identity Server to access backend APIs. Access to the UI is exclusively granted to authorized users, with authentication and authorization handled by a dedicated Authentication/Authorization unit, responsible for controlling user access and logging into the application.
9. Additional public services can offer crucial meteorological data, weather forecasts, maritime information, alerts, and more for visualization within the Dashboard UI.

The connection between the Backend Coordinator and the MQTT system is bidirectional. If any data needs to be transmitted from the application outward, the Backend Coordinator places it in MQTT, within the corresponding queue.

B. Backend Coordinator

The proposed Tool incorporates the following services:

- *MongoDbService*. This background hosting service is responsible for writing all events received from the corresponding Detection Services to the "UAV" and "SAT" collections.

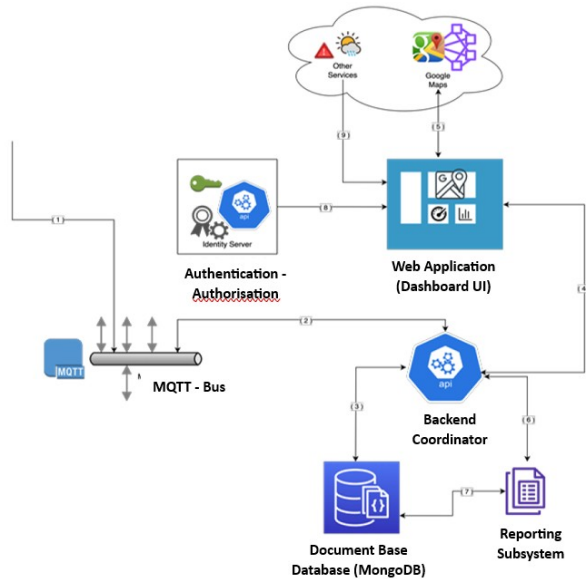


Fig 1. Tool's Architecture Diagram

- *MqttSubscriberService*. This background hosting service is responsible for listening to the MQTT topics and forwarding these messages to the Dashboard UI using WebSocket communication for real-time event presentation. It utilizes the MongoDbService to store these messages in the Document-Based Database (MongoDB) Backend Inventory.
- *ReportingService*. This background hosting service is tasked with generating filters based on criteria selected by end users on the Dashboard. It listens to user requests from the Dashboard for searching and reporting purposes. Upon receiving these requests, it executes queries on the Document-Based Database (MongoDB) Backend Inventory through the MongoDbService of the Backend Coordinator Service. Subsequently, it returns the results to the Dashboard Web App for further visualization and exporting functionalities.

C. Authentication Procedure

Authentication, Authorisation and Audit Logging component, is responsible for intelligently controlling access to UI tools system functions and interfaces (both GUI and REST-API), enforcing policies, and keeping an audit trail of events happening. Based on assigned roles, authenticated users are able to access different UI system functions and interfaces. The audit logging mechanism log several types of information that the system generates during normal execution, such as data changes and actions/commands invoked by the end-users. Structuring the UI web application to support a security token service (Authentication, Authorization and Audit Logging component) leads to the architecture and protocols shown in Figure 2.

This implementation enhances security by ensuring secure authentication, fine-grained authorization, token-based security, secure communication, client credential management, to-

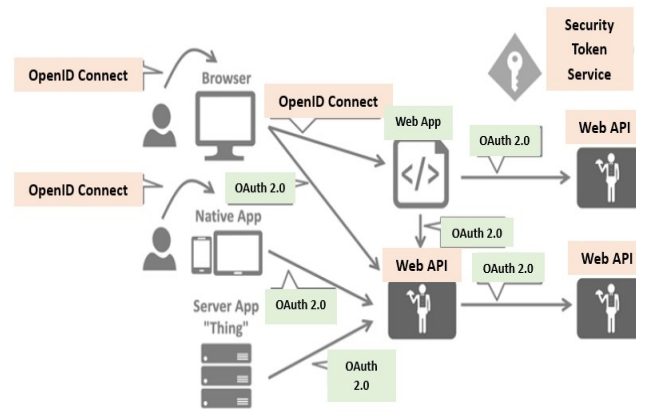


Fig 2. Security Token Architecture and Protocols

ken revocation, and adherence to standardization and best practices. In addition, the Token-based security is a crucial aspect, as it relies on access tokens for secure communication between the Dashboard UI client and the Backend Coordinator backend service. These tokens are short-lived and cryptographically signed, minimizing the risk of unauthorized access and data breaches. Access tokens serve as temporary credentials, granting access to protected resources for a limited duration. Upon expiration, the Dashboard UI client obtains new tokens through the OAuth authentication process, reducing the window of vulnerability and enhancing security. Secure authentication is achieved through the use of OAuth, which replaces the direct sharing of sensitive credentials like usernames and passwords with the issuance of access tokens. These tokens serve as proof of authentication and are included in subsequent API requests to the Backend Coordinator backend service, ensuring secure access to protected resources. The Identity Server issues access tokens during the OAuth authentication process, providing a secure and efficient way to authenticate users.

D. Two-Factor Authenticator of Dashboard UI

To enhance the security access of the Dashboard UI, a Two-Factor Authentication feature using Google Authenticator has been integrated with the existing Identity Server OAuth service. Two-factor authentication (2FA) on web applications works as follows: The user first logs in with their username and password, which represents the first authentication factor - something they know. After successfully entering the username and password, the web application then prompts the user to provide a second form of authentication, such as a one-time code sent to their registered mobile device. This one-time code represents the second authentication factor - something the user has. The user receives the one-time code, typically via SMS or a mobile app, and enters it into the web application to complete the login process. Once the user provides the correct one-time code, they are granted access to the web application.

E. WebSocket of Dashboard UI

In this tool, real-time data presentation between the Backend Coordinator and the Dashboard UI is facilitated through SignalR, a library that implements the WebSocket protocol. This implementation prioritizes security to safeguard sensitive information exchanged in real-time.

IV. MQTT INTEGRATION

Within the Dashboard UI, an MQTT architecture is utilized with a central MQTT broker facilitating a publish-subscribe mechanism for data ingestion from three distinct sources: UAV platform, satellite, and potentially legacy systems. The integration involves the Satellite Component, which is part of the Backend Coordinator system, interacting with the Dashboard UI to facilitate the prediction process based on user-defined areas of interest (Fig.3).

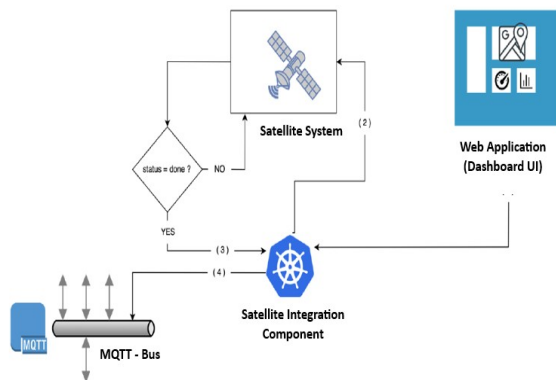


Fig 3. Satellite Component Diagram

As far as, the integration which involves the UAV Component, contains a addition to the Dashboard's internal infrastructure, collaborating with the UAV Platform and UAV Detection System to enhance the detection process (Fig.4).




Fig 4. UAV Integration Component Diagram

V. DEPLOYMENT

The Eclipse Mosquitto MQTT broker and the MongoDB database are key components of the overall system architecture, providing the necessary messaging and data storage capabilities to support the various subsystems and the Dashboard UI. Specifically:

- *MQTT Bus Service*: this involves installing the Eclipse Mosquitto MQTT broker, version 5 or 3.1.1, on the cloud platform, configured to listen on default ports (1883). The initial setup of the broker includes

creating two main topics: "UAV" for use by the UAV detection subsystem and "SAT" for use by the Satellite Component. Both subsystems publish the results of their detection processes to these topics.

- *Document-Based Database (MongoDB) for Backend Inventory*: this entails installing the latest version of MongoDB on the cloud platform, configured to listen on default ports (27017). The initial setup of the database includes a database named "Sunrise" with two collections: "UAV" for storing all detection events from the UAV detection subsystem and "SAT" for storing all detection events from the Satellite Component. These events are later utilized for reporting and historical data visualization on the Dashboard UI.

In addition, in Tool's design, a comprehensive suite of security features has been integrated to enhance the integrity, confidentiality, and reliability of our MQTT communication protocol. These measures collectively fortify the security posture of our system, ensuring data protection and secure communication channels. The next figure shown the integration diagram with MQTT Broker of this tool:




Fig 5. Integration Diagram

Finally, the Integration Process described as follow: The Publishers establish a connection to the MQTT broker and authenticate themselves using credentials. Once connected, publishers can start sending their messages. The MQTT broker receives these Messages containing data and topic information and temporarily stores them. Subsequently, the MQTT broker routes the messages to subscriber based on subscriber's interest. On data integration the subscriber receives the messages that containing relevant data and proceeds as needed. Integration can involve storing data in databases, triggering actions, generating notifications, or updating visualizations.

On Data validation in an MQTT system, a crucial step is to ensure that the incoming data from publishers is accurate, consistent, and conforms to the expected format. This helps prevent erroneous data from being distributed to subscribers and ensures the overall integrity of the system. Data validation in an MQTT system is based on several parameters which are Payload Format Validation, Subject Validation, Data Range and Constraints, Message Size Validation, Protocol Validation, Quality of Service (QoS) and Identification, Flag Validation Preservation, Security Checks and Error Handling.

VI. DASHBOARD MOCKUPS

The SUNIRISE login page presents a secure gateway to access the platform. Users input their credentials – a unique

combination of username and password – to verify their identity. The page's design is intuitive, with fields for entering credentials prominently displayed. Once verified, users gain authorized entry, unlocking the platform's features and personalized /CI related content. In case of forgotten credentials, the page also provides options for password recovery or account assistance. After granting access to the SUNRISE platform, the user navigates to the first page with the GIS Map and relevant layers of the infrastructure of their responsibility. On the right side a list of all the events that have been detected is presented with some fundamental information regarding each specific event.

By clicking the marker on the map or an event on the Events List on the right side, the annotated images from the inspected infrastructure are shown on a pop-up window in the main screen of the GIS (Fig.6). This pop-up window also presents the fundamental info such as event type, time stamp, location, source of inspection etc.

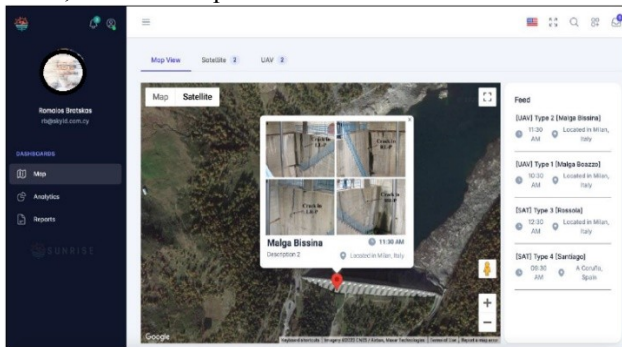


Fig 6. Event Presentation

Finally, for example, we can have the list of the events by source category (Satellite – UAV)(Fig.7).

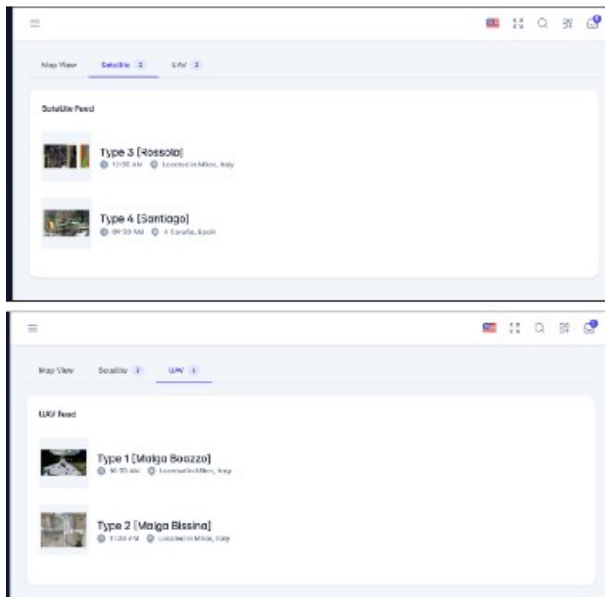


Fig 7. List of the events by source category (Satellite, UAV)

VII. CONCLUSION

The primary goal of the user interface (UI) for remote infrastructure inspection is to provide users with real-time images of critical infrastructure, highlighting areas, components, or points of failure such as damaged parts, structural issues, corrosion, and obstructing vegetation. This is achieved by analyzing image and video feeds and presenting outputs from AI-assisted components into the UI. The design and implementation of the user interface as well as the inspection functionalities provided that required by SUNRISE project, have been based on the following components:

- Two inspection data sources which are a UAV platform and satellite imagery. Both systems through AI technologies can detect the types of problems CI operators are concerned about.
- The interconnection and exploitation of any legacy systems which may be at the CI.
- The visualization of the imports in the form of lists of events. These imports are data that have already been annotated by the corresponding inspection tool.
- The reporting services where statistics about the inspections as well as accessing historical information will be provided – if existing.

ACKNOWLEDGMENT

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This research was funded by SUNRISE H2020, which has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101073821.

REFERENCES

- [1] Jaroshaw, M. (2022). Location Accuracy of a Ground Station based on RSS in the Rice Channel. Proceedings of the 17th Conference on Computer Science and Intelligence Systems, pp. 577-80, dx.doi.org/10.15439/2022F15
- [2] Buczyński, H. Pisarczyk, P. and Cabaj, K. (2022). Resource Partitioning in Phoenix-RTOS for Critical and Noncritical Software for UAV systems. Proceedings of the 17th Conference on Computer Science and Intelligence Systems, pp.605-9, http://dx.doi.org/10.15439/2022F163
- [3] Danilchenko, K. and M. Segal (2021). An Efficient Connected Swarm Deployment via Deep Learning. Proceedings of the 16th Conference on Computer Science and Intelligence Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, D. Ślęzak (eds). ACSIS, Vol. 25, pages 1–7 (2021).http://dx.doi.org/10.15439/2021F001
- [4] Adam, T. and F. Babic (2021). UAV Mission Definition and Implementation for Visual Inspection. Proceedings of the 16th Conference on Computer Science and Intelligence Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, D. Ślęzak (eds). ACSIS, Vol. 25, pages 343–346 (2021) http://dx.doi.org/10.15439/2021F24
- [5] A. Dillon, (2006). User Interface Design. DOI: 10.1002/0470018860.s00054
- [6] A. Panda, J.N. Ramos, Making Critical Infrastructure Resilient: Ensuring Continuity of Service Policy and Regulations in Europe and Central Asia, 2020. www. undrr.org.

- [7] W. Clinton, Presidential Decision Directive 63; The White House: Washington, DC, USA, 1998. Available online: fas.org/irp/offdocs/pdd/pdd-63.htm (accessed on 29 April 2024).
- [8] National Infrastructure Advisory Council (US); T. Noonan; E. Archuleta. The National Infrastructure Advisory Council's Final Report and Recommendations on the Insider Threat to Critical Infrastructures; DHS/NIAC: Washington, DC, USA, 2009.
- [9] L. Coppolino, L. S. D'Antonio, S.V. Giuliano, G. Mazzeo and L. Romano. A framework for Seveso-compliant cyber-physical security testing in sensitive industrial plants. *Comput. Ind.* 2022, 136, 103589. <https://doi.org/10.1016/j.compind.2021.103589>.
- [10] Monstadt, and M. Schmidt, Urban resilience in the making? The governance of critical infrastructures in German cities 56 (11) (2019) 2353–2371, <https://doi.org/10.1177/0042098018808483>
- [11] G.P. Cimellaro, P. Crupi, H.U. Kim, and A. Agrawal, Modeling interdependencies of critical infrastructures after hurricane Sandy, *Int. J. Disaster Risk Reduc.* 38 (2019), 101191, <https://doi.org/10.1016/J.IJDRR.2019.101191>.
- [12] B. Rathnayaka, C. Siriwardana, D.J. Robert, P. Amaratunga, and S. Sujeeva. Improving the resilience of critical infrastructure: Evidence-based insights from a systematic literature review. *International Journal of Disaster Risk Reduction*, 2022, <https://doi.org/10.1016/j.ijdr.2022.103123>.
- [13] OECD, Good Governance for Critical Infrastructure Resilience, OECD, 2019, <https://doi.org/10.1787/02F0E5A0-EN>.
- [14] F. De Felice, I. Baffo, and A. Petrillo. Critical Infrastructures Overview: Past, Present and Future. *Sustainability* 2022, 14, 2233. <https://doi.org/10.3390/su14042233>.
- [15] C. Berger, P. Eichhammer, H.P. Reiser, J. Domaschka, F.J. Hauck, and G. Habiger. A Survey on Resilience in the IoT: Taxonomy, Classification, and Discussion of Resilience Mechanisms. *ACM Comput. Surv.* 2022, 54, 147. <https://doi.org/10.1145/3462513>.
- [16] M. Barnes, K. Brown, J. Carmona, D. Cevasco, M. Collu, C. Crabtree, W. Crowther, S. Djurovic, D. Flynn, P.R. Green, P.R. et al. Technology Drivers in Windfarm Asset Management. *Home Offshore*. 2018. Available online: <https://doi.org/10.17861/20180718> (accessed on 15 May 2024).

Evolving the Enterprise Software Systems Landscape: Towards Anti-Patterns in Smalltalk-to-Java Code Transformation

Marek Bělohoubek and Robert Pergl

Faculty of Information Technology, Czech Technical University in Prague, Prague, Czech Republic

Email: {marek.belohoubek,robert.pergl}@fit.cvut.cz

<http://ccmi.fit.cvut.cz>

Abstract—In the rapidly evolving landscape of enterprise software systems, there is a marked escalation in the proliferation of new technologies, tools, languages, and methodologies daily. These innovations are pivotal not only for the development of new systems but also for the maintenance and augmentation of existing infrastructures. Consequently, it is imperative to devise systems that are responsive to these advancements, fostering the integration of novel tools and methodologies into the current systems. This integration often necessitates mechanisms for transforming source code across diverse programming languages. In the course of developing a transformation tool from Smalltalk to Java, we encountered several code patterns that significantly impede the transformation process. This paper aims to elucidate one such transformation anti-pattern. We provide a comprehensive overview, a formal delineation, illustrations derived from actual code, and propose refactoring strategies for both Smalltalk and Java environments.

I. INTRODUCTION

THE EVOLUTION of software systems within enterprises is a perpetual and intricate process. Each day ushers in a profusion of fresh technologies, tools, programming languages, and methodologies aimed at facilitating the development of novel systems while also helping in the maintenance and expansion of existing ones [1].

The significance of maintenance and expansion cannot be overstated, primarily because as these systems mature, there emerges a necessity to adapt to updated requirements, thereby compelling the integration of new libraries, the adoption of the latest technologies, and, in extreme cases, the complete restructuring of the underlying architecture.

As current research suggests, the frequency of such modifications is increasing [2] it becomes essential to develop new systems that take into account these modifications [3] and create tools and methodologies to help incorporate this design philosophy into existing systems, for example by helping to transform source code between different programming languages.

Over a year ago, we began researching the transformation of Smalltalk code to Java, based on previous research by Engelbrecht [4], with the aim of creating tools and methodologies following these three principles:

- Transformation process should require as little manual input as possible.

- The resulting code should be easy to read and edit manually, to make future expansion and maintenance as simple as possible.
- Tools need to provide option to manually override (sub)result of any step in the process, and allow inclusion of these changes during repeated processing of the input code.

Throughout the process of implementation and subsequent testing on real-world code, we encountered numerous instances where our methods failed to effectively transform certain code snippets.

Upon closer examination of these snippets, we identified recurring code patterns that, while perfectly valid in Smalltalk, lacked direct equivalents in Java suitable for automated translation without compromising the original functionality or excessively cluttering the resulting code.

Internally, we have coined the term "Transformation anti-patterns" to describe these phenomena. In this paper, we aim to elucidate one such anti-pattern by furnishing a comprehensive description, formal definition, real code-based example, and refactoring solutions applicable to both the Smalltalk and Java contexts.

II. LANGUAGES

First, we need to provide the reader with additional information about the languages and design of our tool.

A. Smalltalk

Smalltalk, which emerged in the 1970s as Smalltalk-72, is a dynamically typed programming language that is purely object-oriented. It was first made available to the public with the release of Smalltalk-80 in the 1980s [5].

Smalltalk is distributed in the form of an image along with its own development environment and virtual machine. Consequently, every application developed in Smalltalk necessitates the execution of the associated image.

Various versions of Smalltalk (called dialects) exist, each with its own interpretation of the Smalltalk virtual machine. These include proprietary systems such as Visual Works [6], as well as open-source initiatives like Squeak [7] and Pharo [8].

B. Java

Java, created in the 1990s and launched in 1996 as Java 1.0, is a statically typed, object-oriented, high-level programming language [9].

A crucial component of Java is the Java Virtual Machine (JVM) [10], which facilitates the running of Java applications on various computer architectures, regardless of the platform used. This abstraction shields programmers from platform-specific intricacies.

C. Language differences

There are three main distinctions between Smalltalk and Java that complicate the conversion process to the extent that we refer to it as transformation.

Firstly, they vary in their handling of typing and types. Java follows a static typing approach, requiring all variable types, method arguments, and return values to be defined before compilation. In contrast, Smalltalk is dynamically typed, meaning type checks are performed only during run-time.

The second differentiation lies in the contrast between Smalltalk's meta-classes and Java's class features. In Smalltalk, each object holds a reference to its metaclass, which stores class-specific methods and variables. Conversely, Java uses the static keyword to define class features, which do not pertain to an underlying object and thus cannot be accessed through reference.

Finally, whereas all types in Smalltalk are subclasses of Object, Java includes eight primitive types that solely contain values without any Object-like characteristics (such as methods, initialization, etc.).

These disparities necessitate transformation tools to either deduce all types used in the code for integration into statically typed Java, and/or implement an additional framework to mimic Smalltalk's dynamic behavior.

III. SMALLTALK TO JAVA TRANSFORMATION TOOL

This section consists of two parts: general overview of the transformation process and a more in-depth look at the translation step.

A. Transformation process

Our vision for the tool primarily revolves around facilitating the transition of the business logic from legacy Smalltalk systems into contemporary object-oriented languages.

These systems often grapple with challenges related to maintaining comprehensive and accurate documentation due to extensive years of ongoing maintenance and subsequent development. Given that the tool necessitates parsing the source code of the system for transformation purposes, it can concurrently generate a class model of the system as an ancillary outcome.

Considering that the transformation process will likely entail substantial modifications to the underlying architecture and technologies, particularly concerning the user interface (UI), it becomes imperative to support only partial transformation of the system.

In light of these requirements, our approach involves the development of a suite of tools designed to execute the transformation through the following sequential steps:

- **Scope definition** User defines the transformation scope by specifying which parts (bundles, packages, classes) of the system should be transformed.
- **Type inferring** Transformation tool then uses a real-time type inferrer [11] to obtain all types used in the selected scope.
- **Translation** The scope is then transformed class-by-class, method-by-method, using previously inferred types. The source code is translated from Smalltalk to Java, so the result not only works like the original but also looks as close as possible to the original.
- **Model generation** The transformation tool then produces a model of the transformed scope, in the form of a UML with a transformation profile applied to it.
- **Java generation** Finally, the generation tool uses the transformation model to generate the package structure and Java classes out of it.

B. Translation step

As mentioned previously, our goal is to make a set of tools and methodologies to help programmers migrate applications and systems from Smalltalk to Java with the expectation of further maintenance and development performed within the transformed code.

Therefore, such tools have to be capable of producing human-readable and editable code, which limits usage of certain methods that deal with the discrepancy between dynamically typed Smalltalk and statically typed Java in elegant ways, but produce code that is very hard to maintain and expand.

One such method described in Mr. Engelbrecht's work [4] suggests using a single class named `SmalltalkObject` as an ancestor to all classes in the transformed system.

This class then defines a dummy definition for every method in the original system with all inputs and outputs replaced by `SmalltalkObject` itself. Concrete classes override their own methods with specific implementations.

With this clever use of class-based reflection, this method effectively simulates Smalltalk behaviour in Java, but at the cost of extreme inconvenience for the future expansions of the transformed code, since every class implements or inherits every single method defined in the transformed code.

C. Translation tool

Our implementation of the tool responsible for the translation step treats each method / class definition as a separate code snippet, with the only connection to the rest of the transformed scope being inferred types.

We use the abstract syntax tree generated by Smalltalk compiler to transform methods node-by-node, in essence doing word-by-word literal translation. This is combined with usage of the inferred types to produce being human-like strongly typed code in Java.

To better demonstrate here is an example of simple Smalltalk method:

```
sum: aFirst and: aSecond
  ^ aFirst + aSecond.
```

Assuming that we have previously inferred that `aFirst`, `aSecond` and their sum are all typed `Integer`, we will get the following translation:

```
public Integer sumAnd(Integer aFirst,
  Integer aSecond){
  return aFirst + aSecond;
}
```

As demonstrated in the example above, the translated code uses strong typing and for simpler methods looks like it was written by a human.

Unfortunately during our testing on real-world application we have discovered several code patterns that prevent fully automated translation, either because there is no direct strongly typed equivalent to the original weakly typed code, or because the automatic solution would introduce severe complications to readability and extensibility of the code.

We have named such code structures "Transformation anti-patterns".

IV. TRANSFORMATION ANTI-PATTERNS

We define transformation anti-patterns as code structures that satisfy all of the following conditions:

- The code structure is syntactically and semantically valid in Smalltalk, capable of being compiled and executed without any errors.
- Translating the code structure directly from Smalltalk to Java would yield Java code that either fails to compile or exhibits considerable divergence in runtime behavior when compared to the original Smalltalk implementation.

Virtually all anti-patterns identified thus far originate from the semantic discrepancies between the two languages, frequently stemming from the contrasting principles underlying dynamic and static typing.

This phenomenon can be likened to the challenges inherent in translating between natural human languages. For example, the literal translation of a joke that relies on clever wordplay and phonetic similarity in one language will almost invariably fail to convey its original humour in another language.

Therefore, it is imperative to detect such occurrences during the translation process and notify the user, advising them to review the outcome, and, if possible, offering guidance on resolving any issues that may arise.

Our investigation into the transformation from Smalltalk to Java necessitates the identification of suspicious code structures, their categorisation into anti-patterns, and the provision of multiple strategies for addressing them.

The following section presents an example of one such anti-pattern, identified during experiments on a real-world

Smalltalk system (regrettably, direct examples of the code cannot be provided due to an existing non-disclosure agreement).

While the example anti-pattern isn't the most complicated one, it demonstrates the basic principles behind anti-patterns very well and we have found it occurring quite frequently in the analysed code.

V. ANCESTOR DEFINED VARIABLE

The employment of inheritance and class hierarchy is primarily motivated by the opportunity to establish methods and variables within a high-level class and subsequently allow all its subclasses to inherit these attributes, thereby adhering to the "Don't Repeat Yourself" (DRY) principle.

This practice is prevalent in both Smalltalk and Java, typically manifesting in the form of abstract classes or Java interfaces. However, disparities arise in its utilization between the two languages due to the distinction between dynamic and static typing.

In dynamically typed languages like Smalltalk, programmers can define variables and methods in ancestor classes with a general type and then utilize specialized types in subclasses instead.

While this approach is also feasible in strongly typed languages like Java, it necessitates explicit casting each time the programmer accesses the actual value or return value of a method from the ancestor class, to align with the specialized type in the subclass.

This casting can be implemented either by incorporating casts wherever the variables/methods from the ancestor class are utilized, or by overriding pertinent methods with new implementations that invoke their parent counterparts and then cast the result to the appropriate type.

However, both of these practices are suboptimal: the former clutters the code with type casts, while the latter violates the DRY principle by necessitating the reimplementations of methods.

Alternatively, programmers can leverage Java generics to define problematic types in the ancestor class and defer the specification of concrete types to its descendants. While this approach typically functions effectively, it encounters two limitations.

Firstly, descendants are only considered polymorphic if they all employ identical concrete types. Secondly, the concrete types must share at least one common ancestor or implement the same interface to be usable in both the ancestor class and the concrete classes, assuming the ancestor class doesn't solely define simplistic `get/set` methods.

While employing disparate types in ancestor and descendant classes is generally discouraged, even in dynamically typed languages, ancestor classes following this pattern are typically abstract, serving as "interfaces with partial implementation," either explicitly designated as such or never instantiated within the program itself.

VI. EXAMPLE

Let's suppose following scenario. We are translating three classes:

- Ancestor - which defines an instance variable called property and implements its accessor.
- DescendantString - subclass of Ancestor, initializes its instance variable with empty String.
- DescendantInt - subclass of Ancestor, initializes its instance variable with int.

Here is code snipped with their Smalltalk implementations:

```
Object subclass: #Ancestor
instanceVariableNames: 'property.'

Ancestor>>property
  ^property.

Ancestor subclass: #DescendantString

DescendantString>>initialize
  property := ''.

Ancestor subclass: #DescendantInt.

DescendantInt>>initialize
  property := 0.
```

Transforming the code above directly is certainly possible (see the end of this subsection), but since each of the Ancestors subclasses initializes property with a completely different type, it will be defined as Object.

Therefore, even if the user wants to work directly with one of the subclasses, he will be forced to use type casting for all but the simplest operations with property.

```
public class Ancestor {
    protected Object property;

    public Object getPropertyValue(){
        return property;
    }
}

public class DescendantString
    extends Ancestor{

    public DescendantString(){
        property = "";
    }
}

public class DescendantInt
    extends Ancestor{

    public DescendantInt(){
        property = 0;
    }
}
```

VII. REFACTORING SOLUTION

There are two main solutions for refactoring that coincide with the place where they occur: Ancestor abstraction (in Smalltalk) and Ancestor parameterisation (in Java).

A. Ancestor abstraction

Move all problematic variables and method implementations from the ancestor to its direct subclasses, making both the ancestor and its problematic methods abstract.

```
Ancestor subclass: #DescendantString
instanceVariableNames: 'property.'

DescendantString>>property
  ^property.

DescendantString>>initialize
  property := ''.

Ancestor subclass: #DescendantInt
instanceVariableNames: 'property'.

DescendantInt>>property
  ^property.

DescendantInt>>initialize
  property := 0.
```

The code above shows the changes in Smalltalk code and the translated result is bellow.

```
public abstract class Ancestor {
    public abstract Object getPropertyValue();
}

public class DescendantString
    extends Ancestor{

    protected String property;

    public DescendantString(){
        property = new String();
    }

    public String getPropertyValue(){
        return property;
    }
}

public class DescendantInt
    extends Ancestor{

    protected Integer property;

    public DescendantInt(){
        property = new Integer();
    }

    public Integer getPropertyValue(){
        return property;
    }
}
```

```
}
}
```

This approach breaks the DRY principle, but the refactoring can be done in the source system (before the transformation to the statically typed language). It also keeps the polymorphism intact.

B. Ancestor parameterisation

Change the problematic ancestor into parametric class (for example, in Java: `Ancestor <T> class`) and define the correct concrete type in the generalisation itself (to continue with our previous Java example: `Descendant class extends Ancestor<concreteType>`).

```
public abstract class Ancestor <T>{
    protected T property;

    public T getPropertyValue() {
        return property;
    }
}

public class DescendantString
    extends Ancestor<String>{

    public DescendantString() {
        property = "";
    }
}

public class DescendantInt
    extends Ancestor<Integer>{

    public DescendantInt() {
        property = 0;
    }
}
```

This approach follows the DRY principle, but moves all of the refactoring to the target side. Worse yet, this approach can break the principle of substitution depending on the target system.

For example, in Java only descendants that have their concrete type matching the currently used ancestor will be allowed for substitution (this is due to the Java generics implementation) [12].

Lastly, this solution can be applied only to instance side methods, because static methods cannot be parameterized.

C. Comparison

The main differences between the two proposed solutions lie in the side of refactoring (source or target) and if the solution breaks or follows the DRY principle.

During our experiments, we have found that *Ancestor abstraction* is the most commonly used, as it provides consistency during repeated transformations (by making all the changes in the source) and keeps the principle of substitution intact, both of which greatly outweigh breaking the DRY principle.

However, it is necessary to always keep the context in mind as there are many cases where *Ancestor parameterisation* leads to better results.

VIII. RELATED WORK

The foundational principles behind transformation anti-patterns are not unique to the transition from Smalltalk to Java [13]. Instead, they are commonly observed across transformations from weakly-typed to strongly-typed languages.

Historically, programmers have grappled with challenges arising from the fundamental differences between the source and target programming languages. Examples include the translation of nested routines from Pascal to C [14], addressing structured programming constraints in the translation from Fortran to C [15], and, most critically, issues related to differences in typing systems, such as those encountered in the transition from Python to Java [16].

Furthermore, the development of transpliers—translating compilers—is an essential area of exploration. Some transpliers are specialized for specific languages, like those for C-to-Rust [17]. However, there is also burgeoning research in the realm of multilingual transpliers [18] [19].

This discussion would be incomplete without acknowledging advancements in AI translation tools. ORIGIN-Transcoder [20] employs a neural-based algorithm to reduce the need for manual input in the translation process. Although its applicability to translations from Smalltalk remains unproven, there is significant interest in the potential role of generative AI in the translation process [21].

IX. CONCLUSION

In this article, we have introduced the concept of transformation anti-patterns and shown an example of one of them.

Anti-patterns themselves will not always cause problems, but they represent severe complications for maintenance and future expansions of the translated code.

Based on this analysis, we believe that there is a space for further research into this topic, with many more anti-patterns yet to be discovered.

ACKNOWLEDGMENT

This research was supported by the grant of Czech Technical University in Prague No. SGS23/206/OHK3/3T/18.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Statement on the use of AI AI technologies (Writefull and ChatGPT) were used solely to improve the language of the paper.

REFERENCES

- [1] M. Hilbert and P. López, “The world’s technological capacity to store, communicate, and compute information,” *science*, vol. 332, no. 6025, pp. 60–65, 2011.
- [2] R. Kurzweil, *The Law of Accelerating Returns*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 381–416.

- [3] O. Dvořák and R. Pergl, "Tackling rapid technology changes by applying enterprise engineering theories," *Science of Computer Programming*, vol. 215, p. 102747, 2022.
- [4] R. L. Engelbrecht *et al.*, "Implementing a smalltalk to java translator," Ph.D. dissertation, University of Pretoria, 2006.
- [5] A. C. Kay, *The Early History of Smalltalk*. New York, NY, USA: Association for Computing Machinery, 1996, p. 511–598.
- [6] "Custom Software Application Development Services - Cincom VisualWorks® | Cincom Smalltalk®," 9 2023. [Online]. Available: <https://www.cincomsmalltalk.com/main/products/visualworks/>
- [7] Squeak.org, "Squeak/Smalltalk." [Online]. Available: <https://squeak.org/>
- [8] "Pharo - Welcome to Pharo!" [Online]. Available: <https://pharo.org/>
- [9] I. Cosmina, *An Introduction to Java and Its History*. Berkeley, CA: Apress, 2022, pp. 1–31.
- [10] F. Yellin and T. Lindholm, "The java virtual machine specification," 1996.
- [11] J. Blizničenko and R. Pergl, "Generating UML Models with Inferred Types from Pharo Code," in *International Workshop on Smalltalk Technologies*, Koln, Germany, Aug. 2019. [Online]. Available: <https://hal.science/hal-04053497>
- [12] M. Naftalin and P. Wadler, *Java Generics and Collections: Speed Up the Java Development Process*. O'Reilly Media, 2006. [Online]. Available: <https://books.google.cz/books?id=zaoK0Z2STlkC>
- [13] M. Bělohoubek and R. Pergl, "The state of smalltalk to java transformation: Approaches review," in *World Conference on Information Systems and Technologies*. Springer, 2024, pp. 235–241.
- [14] N. Sundaresan, "Translation of nested pascal routines to c," *ACM Sigplan Notices*, vol. 25, no. 5, pp. 69–81, 1990.
- [15] D. S. Higgins, "A structured fortran translator," *ACM SIGPLAN Notices*, vol. 10, no. 2, pp. 42–48, 1975.
- [16] E. Jin and Y. Sun, "An algorithm-adaptive source code converter to automate the translation from python to java," *JLPEA*, 2020.
- [17] L. Xia, B. Hua, and Z. Peng, "An empirical study of c to rust transpilers," *School of Software Engineering, University of Science and Technology of China, and Suzhou Institute for Advanced Research, University of Science and Technology of China-04/27*, 2023.
- [18] F. Bertolotti, W. Cazzola, and L. Favalli, "* piler: Compilers in search of compilations," *Journal of Systems and Software*, vol. 212, p. 112006, 2024.
- [19] F. Bertolotti, W. Cazzola, D. Ostuni, and C. Castoldi, "When the dragons defeat the knight: Basilisk an architectural pattern for platform and language independent development," *Journal of Systems and Software*, vol. 215, p. 112088, 2024.
- [20] V. Rajathi, M. Harishankar, J. S. DS *et al.*, "Origin-the transcoder," in *2022 1st International Conference on Computational Science and Technology (ICCST)*. IEEE, 2022, pp. 179–182.
- [21] J. D. Weisz, M. Muller, S. I. Ross, F. Martinez, S. Houde, M. Agarwal, K. Talamadupula, and J. T. Richards, "Better together? an evaluation of ai-supported code translation," in *Proceedings of the 27th International Conference on Intelligent User Interfaces*, 2022, pp. 369–391.

Exploring the role of Artificial Intelligence in assessing soft skills

Matteo Ciaschi

National Research Council (CNR)
email: matteo.ciaschi@cnr.it
ORCID: 0009-0009-5119-3563

Marco Barone

University Giustino Fortunato
University of Studies of Foggia
Email: marco.barone@unifg.it

Abstract— Recent research has underscored the pivotal role of soft skills in navigating the complexities of today's workplace dynamics. Soft skills encompass a broad spectrum of attributes, such as effective communication, adept collaboration, nimble adaptability, and profound emotional intelligence, all of which are integral to fostering productive team environments and driving organizational success. Despite their acknowledged importance, quantifying and evaluating soft skills has traditionally been hindered by their inherently subjective nature. However, the emergence of artificial intelligence (AI) technologies has revolutionized the landscape of skill assessment, presenting novel opportunities to address these longstanding challenges. By leveraging AI-powered algorithms, organizations can now analyze vast datasets encompassing various facets of human interaction, enabling a more nuanced and objective evaluation of individuals' soft skill proficiencies. Moreover, AI-driven assessments offer scalability, allowing for the efficient evaluation of large cohorts of employees or candidates. Nonetheless, this intersection of AI and soft skills measurement is not without its obstacles. Ethical considerations surrounding data privacy, algorithmic bias, and the potential for automation-induced job displacement necessitate careful scrutiny and regulation. Furthermore, the dynamic nature of soft skills presents a continuous challenge, as individuals must continually adapt and refine their abilities to meet evolving workplace demands. Despite these challenges, the synergistic relationship between AI and soft skills measurement holds immense promise for the future of talent assessment and development. By embracing AI-driven approaches, organizations can cultivate a workforce equipped with the diverse skill set necessary to thrive in an ever-changing professional landscape.

Index Terms— natural language processing (NLP), artificial intelligence (AI), human resources (HR).

I. INTRODUCTION

IN RECENT years, the integration of Artificial Intelligence (AI) technologies into various aspects of human resource management has garnered significant attention [15]. This trend is particularly pronounced in the realm of assessing soft skills as shown in *Figure 1*, where AI holds the promise of revolutionizing traditional methodologies. One of the key advantages of AI in assessing soft skills lies in its ability to mitigate the shortcomings of conventional approaches. By leveraging machine learning algorithms and natural language processing techniques, AI systems can analyze large volumes of data with unprecedented speed and accuracy [1]. This capability addresses concerns related to biases, inconsistency, and

subjectivity often associated with human-led evaluations. Moreover, AI-driven assessment tools offer scalability, enabling organizations to evaluate soft skills across diverse populations efficiently. Whether in the context of recruitment, performance evaluations, or training programs [16]. AI-powered solutions can streamline the assessment process while maintaining rigor and reliability [17]. However, the integration of AI in soft skills assessment is not without its challenges. Ensuring the fairness and transparency of AI algorithms, for instance, remains a pressing concern. Biases inherent in training data or algorithmic decision-making processes can inadvertently perpetuate existing inequalities or overlook crucial nuances in human behavior [18].

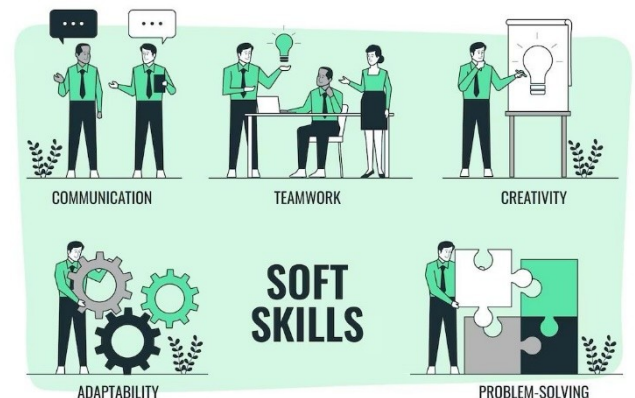


Figure 1 - Example of soft skills

Furthermore, the contextual nature of soft skills poses a unique set of challenges for AI systems. While machine-learning models excel at pattern recognition and prediction, they may struggle to capture the subtleties and nuances of human interaction that characterize soft skills as also demonstrated in *Figure 2*. Nevertheless, the potential benefits of AI in assessing soft skills are substantial. By harnessing the capabilities of AI technologies, organizations can gain deeper insights into the soft skills landscape, identify talent more effectively, and tailor development programs to individual needs. As the field continues to evolve, further research and

innovation will be essential to unlock the full potential of AI in enhancing our understanding and assessment of soft skills. Soft skills often referred to as interpersonal or non-technical skills, play a critical role in professional success across diverse industries. While hard skills are essential for specific tasks, soft skills are equally important for effective communication, teamwork, and leadership. However, quantifying and evaluating soft skills have been traditionally elusive due to their qualitative and context-dependent nature. This paper investigates the evolving landscape of soft skills assessment with the integration of artificial intelligence (AI) technologies.

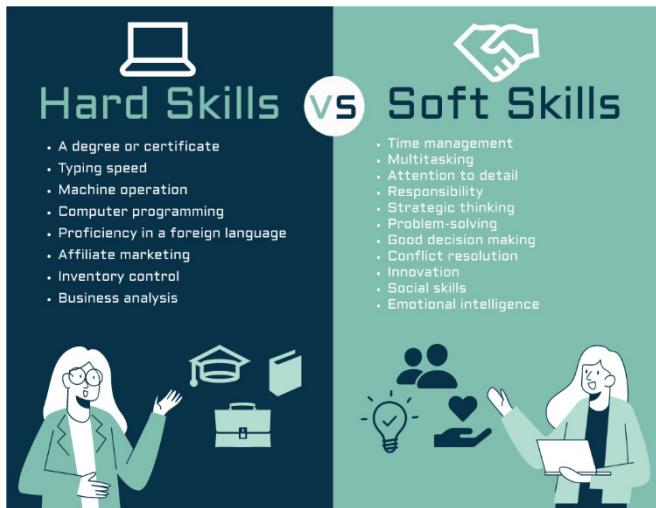


Figure 2 - Hard Skills vs Soft Skills

II. THE IMPORTANCE OF SOFT SKILLS

Soft skills, a diverse set of attributes encompassing communication, empathy, creativity, and problem-solving, have garnered increasing recognition and value from employers. Studies consistently demonstrate that individuals possessing robust soft skills not only thrive in team environments but also exhibit leadership potential and adaptability to change with greater ease. Despite their pivotal role, soft skills have historically received scant attention in both educational curricula and hiring processes. The emergent emphasis on soft skills reflects a fundamental shift in the priorities of modern workplaces. Employers recognize that technical proficiency alone does not suffice in today's dynamic and interconnected business landscape. Instead, the ability to communicate effectively, collaborate harmoniously, and think creatively has become indispensable for fostering innovation, driving productivity, and maintaining competitive advantage [4]. Moreover, the growing complexity of global markets and the rise of digital technologies have intensified the demand for individuals capable of navigating ambiguity and uncertainty. Soft skills, characterized by their flexibility and adaptability, play a critical role in enabling individuals to thrive amidst rapid change and disruption. Yet, despite their demonstrable impact on organizational success, soft skills remain underdeveloped in

many individuals. Traditional education systems, focused primarily on imparting technical knowledge, often neglect the cultivation of essential interpersonal and intrapersonal competencies. Similarly, hiring practices frequently prioritize hard skills over soft skills, overlooking the pivotal role the latter play in fostering collaboration, innovation, and resilience within teams. Recognizing the significance of soft skills is the first step towards addressing this gap. By fostering a culture that values and nurtures these competencies, organizations can unlock the full potential of their workforce and cultivate a dynamic and resilient workplace environment. Embracing this holistic approach to talent development is essential for thriving in an increasingly complex and interconnected world.

III. CHALLENGES IN SOFT SKILLS MEASUREMENT

Measuring soft skills presents a complex endeavor due to their inherent subjectivity and multifaceted nature. Unlike hard skills, which can be objectively assessed through standardized tests or quantifiable performance metrics, soft skills such as communication, teamwork, and emotional intelligence are often more abstract and context-dependent, making their evaluation inherently challenging.

One of the primary obstacles in measuring soft skills is the inadequacy of traditional assessment methods to capture the full spectrum of these skills accurately. Conventional approaches, such as self-assessment surveys or observation-based evaluations, may lack the sensitivity to discern subtle variations in individuals' soft skill proficiency. Consequently, there exists a risk of overestimating or underestimating an individual's soft skills competency, leading to unreliable results. Moreover, there is often a noticeable dissonance between self-reported soft skills and objective assessments conducted by peers or supervisors. Individuals may have biases or lack self-awareness when assessing their soft skills, resulting in discrepancies between perceived and actual proficiency levels. This discordance underscores the importance of incorporating diverse perspectives and utilizing multiple assessment methods to validate soft skills measurement.

Furthermore, cultural, and individual differences add another layer of complexity to soft skills assessment. Cultural norms and expectations can significantly influence how soft skills are expressed and valued, leading to variations in interpretation and evaluation across different contexts. Similarly, individual differences in personality, background, and experiences can impact the manifestation and effectiveness of soft skills, further complicating measurement efforts.

Addressing these challenges requires a multifaceted approach that acknowledges the dynamic and context-dependent nature of soft skills. Innovative assessment methods, such as immersive simulations, real-world scenarios, and behavioral observations, offer promising avenues for capturing the intricacies of soft skills in diverse contexts. Additionally, integrating technology, such as AI and data analytics, can provide valuable insights and enhance the reliability of soft skills assessment tools.

Ultimately, advancing the measurement of soft skills necessitates ongoing collaboration between researchers, educators, employers, and other stakeholders to develop robust evaluation frameworks that are sensitive to individual differences, culturally inclusive, and reflective of real-world demands [2]. By overcoming these challenges, we can better understand, develop, and leverage soft skills to empower individuals and drive success in various personal, academic, and professional domains.

IV. AI-POWERED SOLUTIONS

Recent advancements in AI have applications in many areas including risk management [21], education, communication, healthcare [22], robotics etc. AI tools including natural language processing (NLP), machine learning, and affective computing, offer promising avenues for addressing the challenges of soft skills assessment [20]. NLP algorithms, for instance, have demonstrated remarkable capabilities in analyzing both written and spoken communication, enabling the inference of qualities such as clarity, persuasiveness, and emotional tone with increasing accuracy. These algorithms can sift through vast amounts of text or speech data, extracting meaningful insights that contribute to a more nuanced understanding of an individual's communication skills [3]. Machine learning models, fueled by large datasets, have emerged as powerful tools for identifying patterns in behavior and communication that are indicative of specific soft skills. By analyzing diverse sets of interactions, these models can discern subtle cues and nuances that traditional assessment methods might overlook. Through continuous learning and refinement, machine learning algorithms can adapt to evolving contexts and provide increasingly accurate assessments of individuals' soft skill proficiencies [5].


Affective computing techniques represent another frontier in soft skills assessment, offering the ability to analyze non-verbal cues such as facial expressions, voice intonation, and physiological signals. By leveraging advancements in computer vision and signal processing, affective computing systems can decode emotional states, attitudes, and interpersonal dynamics, shedding light on aspects of emotional intelligence and social competence that are essential for effective communication and collaboration.

By integrating these AI-powered solutions into soft skills assessment frameworks, researchers and practitioners can leverage the vast capabilities of technology to overcome longstanding challenges. These innovations not only enhance the accuracy and reliability of soft skills evaluation but also enable more personalized and adaptive approaches that cater to individual differences and diverse contexts. As AI continues to evolve, the potential for transformative advancements in soft skills assessment becomes increasingly apparent, offering new opportunities to unlock the full potential of individuals and organizations alike [6].

V. APPLICATIONS IN TALENT ASSESSMENT AND DEVELOPMENT

The integration of AI-driven soft skills assessment tools represents a transformative leap forward in talent management practices, offering organizations a wide array of applications across recruitment, employee training, and performance evaluation as presented in *Figure 3*. These innovative tools harness the power of AI to streamline processes, enhance objectivity, and cultivate a workforce equipped with the essential soft skills demanded by today's rapidly evolving business landscape.

Beginning with recruitment, AI-driven screening processes have revolutionized traditional candidate selection methods. By leveraging sophisticated algorithms to analyze vast datasets comprising resumes, cover letters, and online assessments, these tools can swiftly identify candidates whose soft skill profiles align closely with the specific needs and objectives of the organization. The efficiency of automated screening not only accelerates the hiring process but also enables recruiters to focus their efforts on engaging with candidates who demonstrate the requisite communication, collaboration, and emotional intelligence competencies. Furthermore, AI-powered interview platforms equipped with natural language



processing capabilities offer deeper insights into candidates' soft skills by analyzing linguistic nuances, communication styles, and behavioral cues, providing invaluable information

Figure 3 - Applications of Artificial Intelligence

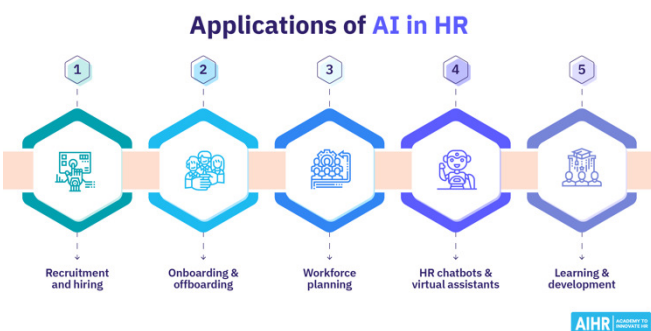
to inform hiring decisions and ensure optimal candidate fit [7]. Moving beyond recruitment, AI continues to play a pivotal role in shaping employee development initiatives. Personalized feedback generated by AI systems provides individuals with granular insights into their soft skill strengths and areas for improvement, empowering them to take ownership of their professional growth journey. Through tailored learning recommendations and resources, employees can embark on targeted skill development paths that align with their unique career aspirations and organizational goals. Additionally, AI-driven coaching platforms leverage real-time data analytics to offer continuous support and guidance, enabling individuals to refine their soft skills in response to evolving workplace challenges and opportunities [8].

In organizational settings, AI-powered analytics serve as indispensable tools for talent management and performance

evaluation. By analyzing team dynamics, communication patterns, and leadership effectiveness, these analytics offer comprehensive insights into the collective soft skill proficiency within teams and across departments. Armed with this invaluable information, organizations can identify opportunities to optimize team performance, allocate resources strategically, and foster a culture of continuous learning and development. Moreover, AI-driven performance evaluations offer objective assessments of employees' soft skill competencies, complementing traditional appraisal methods and ensuring fairness and transparency in talent assessments [9].

The integration of AI-driven soft skills assessment tools represents more than just a technological advancement; it signifies a fundamental shift in how organizations approach talent management and development in the digital age. By harnessing the transformative potential of AI, organizations can unlock the full potential of their workforce, cultivate a culture of excellence, and gain a competitive edge in today's fast-paced and increasingly complex business environment. As AI technologies continue to evolve and mature, the possibilities for enhancing talent management practices are limitless, paving the way for a future where organizations thrive by leveraging the unique strengths and capabilities of their human capital. In recent years, the integration of artificial intelligence (AI) technologies into various domains has revolutionized traditional approaches to problem-solving and decision-making [19]. One area where AI has made significant strides is in Human Resources (HR), where it has transformed talent management practices (**Figure 4**). This paper explores a compelling example of AI application in HR: the use of AI-driven recruitment tools. By leveraging AI algorithms, organizations can streamline the recruitment process, enhance candidate sourcing, and improve decision-making in talent acquisition. This article delves into the benefits, challenges, and implications of AI adoption in HR, shedding light on how these advancements are reshaping the future of workforce management

The advent of artificial intelligence (AI) has brought about a seismic shift in the field of Human Resources (HR), offering innovative solutions to age-old challenges in talent management. With AI's ability to analyze vast amounts of data and



derive actionable insights, organizations are reimagining their

Figure 4- Applications of AI in HR

HR processes to drive efficiency, effectiveness, and inclusivity. One notable application of AI in HR is its utilization in recruitment, where AI-powered tools are revolutionizing how companies identify, attract, and select top talent. This article explores the transformative impact of AI in recruitment, examining its potential to optimize hiring practices and unlock new avenues for talent acquisition.

Traditionally, the recruitment process has been labor-intensive and time-consuming, often fraught with biases and inefficiencies. However, AI-driven recruitment tools offer a paradigm shift in how organizations approach talent acquisition. By leveraging machine learning algorithms, these tools can analyze resumes, assess candidate profiles, and predict job fit with unprecedented accuracy. Furthermore, AI enables organizations to tap into diverse talent pools, mitigating unconscious biases and promoting inclusivity in hiring practices. From automated candidate screening to personalized job recommendations, AI streamlines every stage of the recruitment journey, empowering HR professionals to make data-driven decisions and optimize their hiring strategies.

The integration of AI in recruitment yields a myriad of benefits for organizations. Firstly, AI-driven tools enhance the efficiency of the hiring process by automating repetitive tasks, such as resume screening and candidate matching, freeing up HR professionals to focus on strategic activities. Moreover, AI enables organizations to identify high-potential candidates more effectively, leading to better-quality hires and reduced time-to-fill positions. Additionally, AI-driven recruitment platforms facilitate a seamless candidate experience, providing personalized interactions and timely feedback throughout the application process. By optimizing recruitment practices, AI empowers organizations to build diverse, high-performing teams that drive innovation and competitiveness in the marketplace. Despite its transformative potential, the adoption of AI in recruitment is not without challenges. Ethical considerations, such as data privacy and algorithmic bias, require careful attention to ensure fairness and transparency in decision-making. Moreover, the reliance on AI-driven tools may raise concerns about job displacement and the humanization of the hiring process. To address these challenges, organizations must prioritize ethical AI governance, invest in employee upskilling, and foster a culture of transparency and trust in AI-driven recruitment practices.

Looking ahead, the integration of AI in recruitment heralds a new era in talent acquisition, characterized by data-driven decision-making, enhanced candidate experiences, and greater workforce diversity. As AI continues to evolve, HR professionals must adapt their practices to harness the full potential of these technologies. By embracing AI-driven recruitment tools, organizations can gain a competitive edge in attracting and retaining top talent, positioning themselves for success in the digital age [10].

The application of AI in recruitment represents a watershed moment in the evolution of HR practices. By leveraging AI-driven tools, organizations can optimize their recruitment processes, improve decision-making, and foster a more inclusive

and diverse workforce. However, to realize the full benefits of AI in recruitment, organizations must navigate ethical considerations, address potential biases, and invest in employee development. Ultimately, AI holds the promise of transforming talent acquisition, enabling organizations to build agile, future-ready teams that drive innovation and sustainable growth [11].

VI. DISCUSSION

The integration of artificial intelligence (AI) into the assessment of soft skills has become a focal point in both research and practice, offering innovative solutions to longstanding challenges in talent evaluation. However, this advancement also brings to the forefront a host of ethical considerations that demand scrutiny. Chief among these concerns are issues related to privacy, fairness, and algorithmic bias. In the pursuit of capturing nuanced aspects of human behavior, AI-driven soft skills assessment often relies on the collection and analysis of sensitive personal data, ranging from verbal cues and speech patterns to non-verbal cues such as facial expressions and body language. Such data collection practices necessitate a robust framework grounded in transparency and consent to safeguard individuals' privacy rights and ensure their autonomy in the process. Moreover, the inherently complex nature of human behavior poses significant challenges in developing AI algorithms that can accurately interpret and evaluate soft skills without introducing biases. The risk of algorithmic bias, whereby AI systems unintentionally discriminate against certain individuals or groups, underscores the importance of incorporating diversity and inclusivity considerations into the design and implementation of these technologies. To address these ethical concerns, stakeholders must prioritize proactive measures aimed at mitigating bias and promoting fairness in AI-driven soft skills assessment. This includes adopting strategies to diversify training data, conducting regular audits of algorithmic decision-making processes, and implementing mechanisms for ongoing monitoring and evaluation. Despite these challenges, the potential benefits of AI in advancing soft skills assessment are substantial. By harnessing the power of machine learning and natural language processing techniques, AI systems can offer insights into individuals' interpersonal communication, collaboration, adaptability, and other critical soft skills with unprecedented accuracy and granularity. Looking ahead, future research endeavors should focus on refining AI models through continuous learning mechanisms that enable adaptation to evolving patterns of human behavior. Additionally, the integration of multimodal data sources, such as combining textual and visual information, holds promise for enhancing the comprehensiveness and reliability of soft skills evaluations [12].

Collaboration among interdisciplinary teams comprising AI researchers, psychologists, educators, and industry stakeholders is paramount in driving innovation and maximizing the potential impact of AI on workforce development. By fostering an ecosystem of knowledge exchange and collaboration,

we can collectively address the complex challenges and opportunities inherent in the intersection of AI and soft skills assessment. In the rapidly evolving landscape of the modern workplace, the demand for soft skills continues to grow, fueled by the increasing emphasis on teamwork, creativity, and adaptability. In this context, AI represents not only a tool for improving efficiency and objectivity in talent management but also a catalyst for promoting a culture of continuous learning and development. By leveraging AI technologies thoughtfully and ethically, organizations can unlock new possibilities for nurturing talent, driving innovation, and fostering inclusive and thriving work environments.

VII. CONCLUSION

In conclusion, the symbiosis between AI and soft skills measurement heralds a new era in talent assessment and development. As the demand for soft skills continues to rise in the rapidly evolving workplace landscape, AI presents unprecedented opportunities for transforming how these skills are measured and evaluated. By harnessing the capabilities of AI technologies, organizations can delve deeper into individual and team competencies, transcending the limitations of traditional assessment methods. AI-driven analyses offer nuanced insights into the subtle nuances of human interaction, providing stakeholders with a richer understanding of employees' strengths and areas for improvement [13].

Moreover, AI-powered assessments enable more informed decision-making in talent management, facilitating the alignment of skills with organizational objectives and the strategic deployment of human capital. By identifying and nurturing talent with the requisite soft skills, organizations can cultivate high-performing teams capable of driving innovation and adaptability in today's dynamic business environment.

Furthermore, the integration of AI in soft skills measurement promotes a culture of continuous learning and development. By providing individuals with personalized feedback and targeted interventions, AI-driven platforms empower employees to refine their soft skills iteratively, fostering professional growth and resilience. This emphasis on lifelong learning not only enhances individual performance but also contributes to the overall agility and competitiveness of the organization. [14].

However, it is essential to approach the deployment of AI in soft skills measurement with caution and mindfulness. Ethical considerations, including privacy concerns and algorithmic bias, must be carefully addressed to ensure fairness and transparency in the assessment process. Additionally, organizations must prioritize the upskilling and reskilling of employees to navigate the evolving technological landscape and mitigate the risk of job displacement.

In essence, the convergence of AI and soft skills measurement represents a paradigm shift in how we perceive and cultivate talent in the modern workplace. By embracing AI-driven approaches, organizations can unlock the full potential of their workforce, driving sustainable growth and competitive advantage in an increasingly complex and interconnected world.

REFERENCES

In the case of the bibliography, a thematic order has been followed, with the first six references focusing on pedagogy, multimedia education, and e learning, followed by more specific references on soft skills, emotional intelligence, leadership, affective computing, and performance evaluation.

- [1] Mitchell D. (2008). *What really Works in Special and Inclusive Education*. London: Routledge. (Cit. in Lucio Cottini (2017). *Pedagogia speciale e didattica per l'inclusione*. Roma: Carocci Editore).
- [2] Ranieri, M. (2010). *La Media Literacy nei documenti dell'Unione Europea. Studi e ricerche MED Media Education*.
- [3] Rivoltella, P. C. (2005). Multimedia training, media education, and cooperative learning: new professional scenarios for educators. In A. Ascenzi, M. Corsi (Eds.), *Educator/Trainer Profession. New educational needs and new pedagogical professionalism* (pp. 3-23). Milan: Vita e Pensiero.
- [4] Rivoltella, P. C. (2006). Teacher, mentor, tutor. A framework for reflection on e-learning professionalism. In P. Crispiani, P. G. Rossi (Eds.), *E-Learning. Training, models, proposals* (pp. 55-74). Rome: Armando.
- [5] Rivoltella, P. C. (2017). *Community Technologies*. Brescia: ELS La Scuola.
- [6] Rivoltella, P. C. (2020). *New Alphabets. Education and Cultures in the Post-Media Society*. Brescia: Scholé - Morcelliana.
- [7] Kraiger, K., Ford, J. K., & Salas, E. (1993). Application of cognitive, skill-based, and affective theories of learning outcomes to new methods of training evaluation. *Journal of Applied Psychology*, 78(2), 311-328.
- [8] Salovey, P., & Mayer, J. D. (1990). Emotional intelligence. *Imagination, Cognition and Personality*, 9(3), 185-211.
- [9] Zeidner, M., Roberts, R. D., & Matthews, G. (2008). *The science of emotional intelligence: Knowns and unknowns*. Oxford University Press.
- [10] Goleman, D., Boyatzis, R., & McKee, A. (2013). *Primal leadership: Realizing the power of emotional intelligence*. Harvard Business Press.
- [11] Mayer, J. D., Salovey, P., & Caruso, D. R. (2002). *Mayer-Salovey-Caruso Emotional Intelligence Test (MSCEIT) user's manual*. Toronto, Canada: MHS Publishers.
- [12] Jackson, G. S. (2016). *The art of empathetic leadership: How leaders can drive engagement and build trust*. Palgrave Macmillan.
- [13] Lee, J., Kim, J., Kim, J., Kim, S. J., & Kim, G. (2020). A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 59, 110-125.
- [14] Sackett, P. R., Walmsley, P. T., Lievens, M. L., & Highhouse, M. R. (2017). *Assessment centers and the prediction of managerial performance*. Walter de Gruyter.
- [15] M Naeem, STH Rizvi, A Coronato; A gentle introduction to reinforcement learning and its application in different fields, IEEE access 8, 209320-209344 DOI 10.1109/ACCESS.2020.3038605
- [16] M Naeem, A Coronato, G Paragliola; Adaptive treatment assisting system for patients using machine learning; 2019 sixth international conference on social networks analysis, management DOI 10.1109/SNAMS.2019.8931857
- [17] Naeem, Muddasar and Coronato, Antonio; An AI-empowered home-infrastructure to minimize medication errors, Journal of Sensor and Actuator Networks, V-11.1, P-13, 2022 <https://doi.org/10.3390/jsan11010013>
- [18] A semantic context service for smart offices A Coronato, G De Pietro, M Esposito 2006 International Conference on Hybrid Information Technology 2, 391-399
- [19] A reinforcement learning based intelligent system for the healthcare treatment assistance of patients with disabilities A Coronato, M Naeem International Symposium on Pervasive Systems, Algorithms and Networks, 15-28 DOI 10.1007/978-3-030-30143-9_2

Spoken Language Corpora Augmentation with Domain-Specific Voice-Cloned Speech

Mateusz Czyżnikiewicz, Łukasz Bondaruk,
Jakub Kubiak, Adam Wiacek, Łukasz Degórski
Samsung R&D Institute Poland
Plac Europejski 1
00-844 Warszawa, Poland
Email: {m.czyznikiew,l.bondaruk,j.kubiak3,
a.wiacek2,l.degorski}@samsung.com

Marek Kubis, Paweł Skórzewski
0000-0002-2016-2598
0000-0002-5056-2808
Adam Mickiewicz University, Poland
Faculty of Mathematics and Computer Science
ul. Uniwersytetu Poznańskiego 4
61-614 Poznań, Poland
Email: {mkubis, pawel.skorzewski}@amu.edu.pl

Abstract—In this paper we study the impact of augmenting spoken language corpora with domain-specific synthetic samples for the purpose of training a speech recognition system. Using both a conventional neural TTS system and a zero-shot one with voice cloning ability we generate speech corpora that vary in the number of voices. We compare speech recognition models trained with addition of different amounts of synthetic data generated using these two methods with a baseline model trained solely on voice recordings. We show that while the quality of voice-cloned dataset is lower, its increased multivoiceity makes it much more effective than the one with only a few voices synthesized with the use of a conventional neural TTS system. Furthermore, our experiments indicate that using low variability synthetic speech quickly leads to saturation in the quality of the ASR whereas high variability speech provides improvement even when increasing total amount of data used for training by 30%.

I. INTRODUCTION

WITH THE development of better TTS systems in recent years, there has been an increasing number of research papers on using synthesized data for ASR training [1], [2], [3]. One could argue that, if synthesized samples covered a more diverse set of voice characteristics, even with decrease in speech quality, the data could be used more effectively for training ASR. Conventional neural TTS systems [4], like Tacotron2 [5] or FastSpeech [6], require large amount of high-quality paired text and speech data, which is not available for most languages, especially for multiple voices. Because of that, we cannot use them to produce output with more than a few to a dozen of voices, even for otherwise high-resource languages like German [4]. Recent advancements in speech synthesis brought zero-shot models that use neural codec encoding instead of mel-spectrogram speech representation [7], [8], [9]. Thanks to their zero-shot voice cloning ability, they are able to generate high quality audio with any person's voice, having just a few seconds recording of it. This allows for generating synthetic corpora with hundreds of voices.

Our work examines the usefulness of having a synthetic corpora with a diverse set of voices. For comparison, we

This research was partially funded by the CAIMAC: *Conversational AI Multilingual Augmentation and Compression* project, a cooperation between Adam Mickiewicz University and Samsung Electronics Poland.

employ a zero-shot TTS and a conventional neural TTS to produce a domain-specific synthetic dataset with high and low number of speakers, respectively. We select a virtual assistant (VA) domain as our experiment target. Then, we examine the usefulness of both synthetic datasets in improving the ASR model's performance. We show that the high voice diversity of generated data makes it much more effective. Furthermore, our results indicate that the potential for using synthesized data to improve the ASR performance is limited by variability of the speech produced by a conventional neural TTS system.

II. RELATED WORK

Prior work has shown that using text-to-speech data can improve ASR performance. Rossenbach et al. [3] examined the impact of synthetic data for various ASR architectures. They showed that using TTS data pre-processing techniques can increase the robustness of ASR training. They reported 38% relative improvement after adding synthetic data to the attention encoder-decoder ASR system.

The addition of synthetic data can play an important role in a low-resource setting. Bartelds et al. [10] showed that adding synthetic data to the ASR training on such languages like Besemah and Nasal reduced relative WER up to 25.5%.

In some situations, all that is needed to build an ASR is a text corpus. Rossenbach et al. [11] demonstrated this strategy. They achieved relative improvement of up to 33% in WER over the baseline with data augmentation in a low-resource setting.

Another use for synthetic data can be to improve the recognition of out-of-vocabulary (OOV) words [12]. OOV is a prevalent issue encountered by real-world virtual assistants that must adapt to the ever-evolving environment. Augmentation using TTS-generated data for these specific OOV words can positively affect the robustness of the ASR model without significant degradation on the general dataset.

Kubis et al. [13] use synthesized data to study the impact of speech recognition errors on the performance of natural language understanding models. In [14] text-to-speech models are used in conjunction with an automatic speech recognition system to produce a dataset for improving the robustness of

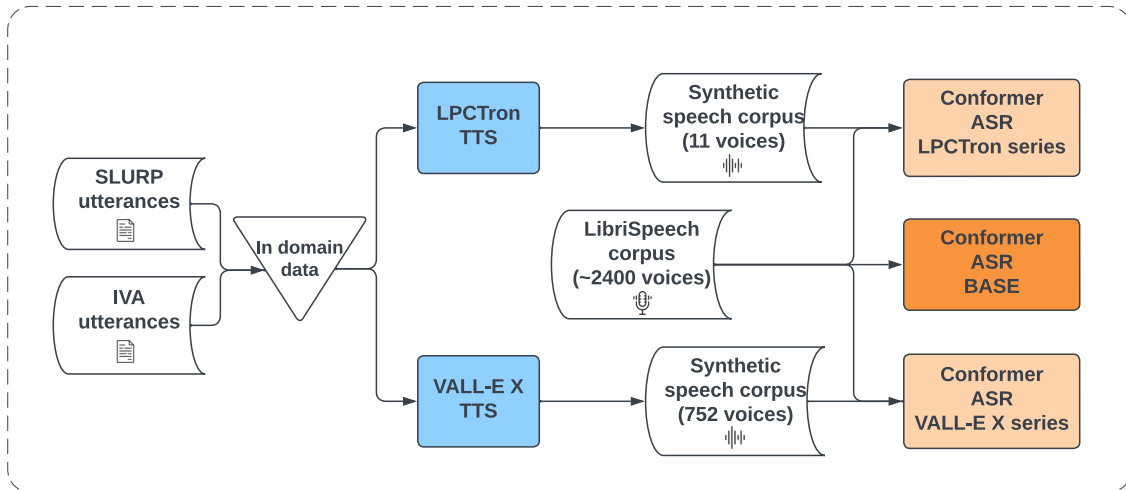


Fig. 1. Experimental workflow.

natural language understanding models to speech recognition errors.

Furthermore, synthetic data might be useful in ASR personalization [15]. The aforementioned study shows high effectiveness in ASR personalization using synthetic data, in particular when there are few recordings of a speaker in the dataset.

Previous works also addressed the problem of imperfections in data produced by TTS. Synthetic data differs from the real one in terms of naturalness and because of the presence of artifacts. Hu et al. [16] proposed two techniques for ASR training to alleviate the issues arising from the problems mentioned above. They observed up to 13% relative error reduction in ASR task.

The authors of VoiceBox [17] investigate the performance of ASR models trained on real and synthetic data. For training the ASR model on real data they use LibriSpeech 100h and 960h datasets. The synthetic data are generated from the texts collected in the LibriSpeech training set. The evaluation is performed with respect to *test-clean* and *test-other* subsets of LibriSpeech which do not contain conversational speech. Le et al. [17] show that their best performing TTS models lead to the absolute WER increase of 0.4% on *test-clean* and 1.7% on *test-other*, if compared to the models trained on real data. Contrary to [17], we investigate the impact of using voice-cloned speech on domain-specific adaptation of ASR in the conversational setting and use for this purpose datasets that contain conversational speech (SLURP and IVA).

III. DATA

Measuring the impact of synthesized data on the performance of the ASR model requires careful selection of speech resources to be used for training and evaluation. We decided to use LibriSpeech [18] as a resource for training baseline ASR model and as a target corpus for augmentation. LibriSpeech is a corpus of approximately 1,000 hours of read English speech, recorded by more than 2,400 speakers. It is derived

from the LibriVox project, which features audiobooks read by volunteers.

For training speech synthesizers we used LJ Speech Dataset [19] and Hi-Fi TTS Dataset [20]. LJSpeech is a dataset of about 24 hours of audio from a single speaker reading book passages, specifically from Project Gutenberg. Hi-Fi TTS Dataset is also based on Project Gutenberg texts and LibriVox audiobooks and contains about 292 hours of speech from 10 speakers with at least 17 hours per speaker. Both of these datasets were designed for training models for speech-based applications, with the main focus on speech synthesis.

We also utilize open-sourced VALL-E X model¹ that was trained on LibriTTS [21], AISHELL-1 [22], AISHELL-3 [23] and Japanese subset of CommonVoice dataset [24]. The authors also used some self-gathered data that was not described. In total they used about 704 hours of speech for English, 598 hours for Chinese and 437 hours for Japanese.

We evaluate ASR models using three general-purpose and two domain-specific ASR datasets. The general-purpose datasets include two test splits of LibriSpeech, *test-clean* and *test-other*. The *test-clean* split has higher quality of samples compared to *test-other* [18]. As a third general-purpose dataset, we use the test split of FLEURS [25] which provides natural speech recordings for many languages, out of which we use an English subset only.

As for the testsets in the domain of virtual assistants, we chose to use the test split of SLURP [26] and our internal virtual assistant (IVA) dataset. The SLURP testset has 13078 recordings totalling 10.3 hours of audio, while the IVA dataset contains 14094 recordings and 12.5 hours of speech. IVA has a broader set of domains and intents (55 and 223 respectively) compared to SLURP (18 and 94). Table I describes the language resources used for evaluation.

For prompting VALL-E X, we randomly chose one record-

¹<https://github.com/Plachtaa/VALL-E-X>

TABLE I
RESOURCES USED FOR EVALUATION.

Dataset	Samples	Hours	Speakers
LS-clean	2620	5.4	40
LS-other	2939	5.1	33
FLEURS	647	1.8	—
SLURP	13078	10.3	—
IVA	14094	12.5	—

ing for each of the speakers. As sources of prompts we used LibriSpeech, HiFi TTS Dataset and LJ Speech Dataset described above and VCTK dataset [27] which contains high quality speech data recorded by 110 English speakers.

IV. MODELS

A. Speech Recognition

For our experiments we chose the Conformer on-device ASR model [28]. It is based on a RNN-Transducer architecture and has been commercialized on edge devices, which proves its high quality. This makes it a compelling target for our experiments on improving the ASR performance.

The model provides real time ASR performance on edge devices. Although the authors used two pass model for better quality, we limited ourselves to the first pass. Our main goal was to observe the difference between both augmentation approaches so we did not find improving ASR by ensembling relevant. In our single pass approach the transcription network encodes acoustic features of speech, while the predictor network acts as language model and tries to predict the next token based on the previous ones. These two, the acoustic features and language features are joined together in the joint network that outputs the final label.

B. Speech Synthesis

As a conventional neural approach to speech synthesis we decided to use a two-stage end-to-end TTS, consisting of an acoustic model mapping phonetic labels to acoustic features and a vocoder mapping these features to audio samples.

The set of phonetic labels contained symbols for phonemes, word delimiters and end of sentence marks (affirmative sentences, questions and exclamations). Acoustic features were derived from F0 (interpolated in unvoiced regions), mel-spectra and band-aperiodicity in a manner of the WORLD vocoder [29]. We utilized vocoder architecture that follows LPCNet [30] and an acoustic model based on Tacotron and [31] Tacotron2 [5], as described in [32]. For simplicity, later we refer to this system as a whole by the name LPCTron.

C. Voice Cloning

VALL-E X [8] is a zero-shot TTS that offers state of the art quality of cloning a sample voice, having only a 3-second recording of it. Instead of regarding speech synthesis as a continuous regression task, it adopts conditional language modelling approach, where the synthesis is conditioned on the input text and audio. It also ceases to use mel-spectrogram in favor of acoustic tokens that are generated by neural codec LM.

The output speech is modeled at two stages with a total of 8 quantizers. In the first stage, the autoregressive language model generates codec codes of the first quantizer. During the second stage, the non-autoregressive language model generates codes for the rest of the quantizers, it is conditioned not on previously generated tokens but on all the tokens from previous quantizers. This makes the second stage much faster, because codes from previous quantizers are known at the start. The intention is that each next quantizer encodes the details that were not captured by previous ones.

The reason that VALL-E X is useful for our task is that it has in-context learning ability, which means that it can synthesize high-quality output on previously unseen inputs. While conventional neural TTS systems needed fine-tuning for unseen speakers, VALL-E X does not.

Open-source VALL-E X implementation follows the original paper [7] and uses G2P tool for converting the input sentence to phonemes and EnCodec [33] as a neural codec.

V. EXPERIMENTS

The goal of our study is to investigate how does the multivoiceity of synthesized, domain-specific training data impact the performance of the resulting ASR model. For this purpose we conduct experiments with ASR models trained on speech recordings, speech recordings combined with data synthesized with LPCTron and speech recordings combined with data synthesized with VALL-E X.

For synthesis, we created a text corpus consisting of 129,000 user commands directed to a task-oriented virtual assistant which includes 81,500 utterances from our internal dataset, and 47,500 utterances obtained in the process of augmenting the training split of the SLURP dataset.

The augmentation employed to enrich SLURP consisted of two steps. First, we used RoBERTa [34] and BART [35] models to randomly substitute words in the user commands with their counterparts supplied by the language models. Second, the sentences were transcribed from English to French, German, Italian and Spanish and backwards with the use of OPUS-MT models [36].

The text corpus was split into 3 equal parts and synthesized using both LPCTron and VALL-E X. For LPCTron we selected voices randomly from 11 available options and for VALL-E X from 752. The audio prompts for VALL-E X were collected from 4 datasets in a manner described in section III. The first part of the text corpus was synthesized with 2 voices per sentence, second part with 3 voices and the last part with 10 voices. This way we obtained three sets of 40 hours, 60 hours and 200 hours of synthesized speech. We combined these sets into splits: 40 hours, 60 hours, 100 hours 200 hours and 300 hours, which were later utilized for experiments.

We used 960 hour subset of LibriSpeech corpus for training along with splits of synthetic data. The L_{xxx} models combine LibriSpeech recordings with LPCTron synthesized dataset with xxx hours, e.g. L_{060} used 60 hours split mentioned above. Analogically, the V_{xxx} models combine LibriSpeech data with xxx hours of spoken commands generated with the use of

TABLE II
LIBRISPEECH 960H ASR MODELS WER.

Dataset	BASE	L040	L060	L100	L200	L300	V040	V060	V100	V200	V300
LS-clean	8.08	7.88	7.62	7.91	8.07	7.80	7.97	9.60	10.74	8.29	8.10
LS-other	20.57	19.84	20.17	20.23	20.51	20.58	20.47	21.43	22.17	20.79	20.75
FLEURS	34.31	34.90	34.02	34.04	34.83	34.44	33.39	33.72	36.28	35.03	33.24
SLURP	74.89	70.02	69.22	68.37	69.56	68.83	66.67	64.64	65.56	63.39	62.58
IVA	75.14	66.82	64.75	62.13	64.09	62.54	50.62	54.01	52.91	47.82	44.61




Fig. 2. WER obtained on SLURP and IVA.

VALL-E X model. The *BASE* model is a baseline trained on LibriSpeech 960h without addition of synthetic data.

The results presented in Table II indicate significant improvement in performance of the augmented models on domain-specific testsets (SLURP and IVA). We can also observe no significant performance drop on general-purpose test sets (LS-clean, LS-other and FLEURS) meaning that ASR models maintained generalization capability. The *V300* performs the best out of all trained models and results in absolute WER reduction, with regard to the *BASE*, of $30.53pp$ and $12.31pp$ in comparison to $12.60pp$ and $6.06pp$ obtained by *L300* on the IVA dataset and SLURP, respectively.

To investigate how the amount of synthetic data used for training impacts the ASR, we compared WER obtained using different data splits of IVA and SLURP. As shown in Figure 2, models trained with addition of VALL-E X data outperform their counterparts augmented with LPCTron data. There is also a noticeable improvement in WER with addition of more

voice-cloned data, whereas the results plateau for models trained with the usage of LPCTron data.

To verify the quality of the audio data produced by VALL-E X and LPCTron we used Whisper [37] ASR model. We computed WER on the subset of 40 hours data. We got 37.55% and 20.38% WER on VALL-E X and LPCTron datasets, respectively.

VI. DISCUSSION

The choice of LPCTron as the baseline for conducting experiments can be questioned as there are several other more recent, conventional neural TTS models that can be used for the task. However, when comparing ratio between MOS for synthesized speech and MOS measured for ground truth across different architectures the results for LPCTron [32] 93% (= 4.2/4.5) are on par with 89% (= 3.83/4.3) achieved for FastSpeech2 [38], 98% (= 4.36/4.45) for HiFiGAN [39] and 93% (= 3.961/4.274) for WaveGlow [40]. Taking into account that HiFiGAN and WaveGlow are vocoders, not the full TTS systems, only FastSpeech2 would be a direct replacement for LPCTron in our experimental setting. Still, FastSpeech2 model presents similar quality to Tacotron2-based TTS models as shown in [38]. Furthermore, as we reported in Section V, the transcriptions of the audio samples produced by LPCTron obtained with the use of Whisper [37] had significantly lower WER than their VALL-E X counterparts. This shows that the quality of generated speech was higher in the case of LPCTron making our study sound, even if the LPCTron model is outperformed by some other conventional neural TTS model that can be potentially used as a baseline for experiments.

Taking into consideration that the compared TTS models are trained in a different manner with VALL-E X being trained for zero-shot (voice cloning) synthesis and LPCTron being trained for a conventional synthesis, there are differences in the model architecture that we cannot control in the experimental setting. However, it should be noted that although VALL-E X is a decoder-only model and Tacotron is an encoder-decoder model both of them are autoregressive, thus we do not consider the differences in the architecture to have a significant impact on the results.

Before VALL-E X, other approaches to zero-shot voice-cloning speech synthesis were considered. They were mainly based on providing the acoustic model with speaker embeddings extracted from speech sample with speaker verification models [41]. This approach still relies on the availability of high quality data for multiple speakers to train acoustic model to utilize speaker embedding space properly. On the

other hand, conditional language modelling approach allows for utilizing lower quality data which makes it more suitable to our study.

VII. CONCLUSIONS

In this study we investigated the efficacy of using voice-cloned speech for augmenting spoken language with the goal of improving the performance of an ASR system. In this setting, we compared a baseline dataset that contains solely voice recordings, the dataset with addition of voice-cloned samples and the dataset expanded with samples synthesized by a conventional neural TTS system.

The conducted experiments show that the use of voice cloning to generate data with multiple voices and pronunciations improves the ASR performance significantly, compared to data from a conventional TTS speaking in just one or a few voices. The lower quality of voice-cloned speech, showed in terms of intelligibility, does not prevent the mentioned improvement.

We also showed that improvements gained by adding more synthetic data to the speech corpus plateau quickly for data generated using conventional neural TTS, but adding even 300 hours of synthetic speech generated using VALL-E X does not seem to saturate the results of ASR model.



One avenue for further research is to investigate upper limits of augmenting speech corpora using voice-cloned samples. Other dimension worth experimenting with is voice characteristics variability and its impact on the ASR results. There is also noticeable gap in quality of synthesized speech in terms of intelligibility between conventional neural TTS and LM-based TTS which should be decreased.

REFERENCES

- [1] A. Fazel, W. Yang, Y. Liu, R. Barra-Chicote, Y. Meng, R. Maas, and J. Droppo, "Synthasr: Unlocking synthetic data for speech recognition," 2021.
- [2] S. Ueno, M. Mimura, S. Sakai, and T. Kawahara, "Data augmentation for asr using tts via a discrete representation," in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2021, pp. 68–75.
- [3] N. Rossenbach, M. Zeineldeen, B. Hilmes, R. Schlüter, and H. Ney, "Comparing the benefit of synthetic training data for various automatic speech recognition architectures," in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2021, pp. 788–795.
- [4] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, "A survey on neural speech synthesis," 2021.
- [5] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomvrgianakis, and Y. Wu, "Natural TTS synthesis by conditioning Wavenet on MEL spectrogram predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 4779–4783.
- [6] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech: Fast, robust and controllable text to speech," 2019.
- [7] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," 2023.
- [8] Z. Zhang, L. Zhou, C. Wang *et al.*, "Speak foreign languages with your own voice: Cross-lingual neural codec language modeling," 2023.
- [9] K. Shen, Z. Ju, X. Tan, Y. Liu, Y. Leng, L. He, T. Qin, S. Zhao, and J. Bian, "Naturalspeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers," 2023.
- [10] M. Bartelds, N. San, B. McDonnell *et al.*, "Making more of little data: Improving low-resource automatic speech recognition using data augmentation," in *Proc. Annual Meeting of the Association for Computational Linguistics*, 2023, pp. 715–729.
- [11] N. Rossenbach, A. Zeyer, R. Schlüter, and H. Ney, "Generating synthetic audio data for attention-based speech recognition systems," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020, pp. 7069–7073.
- [12] X. Zheng, Y. Liu, D. Gunceler, and D. Willett, "Using synthetic audio to improve the recognition of out-of-vocabulary words in end-to-end asr systems," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021, pp. 5674–5678.
- [13] M. Kubis, P. Skórzewski, M. Sowański, and T. Ziętkiewicz, "Back Transcription as a Method for Evaluating Robustness of Natural Language Understanding Models to Speech Recognition Errors," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, December 2023, pp. 11 824–11 835.
- [14] —, "Center for Artificial Intelligence Challenge on Conversational AI Correctness," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Słezak, Eds., vol. 35. IEEE, 2023, pp. 1319–1324.
- [15] K. Yang, T.-Y. Hu, J.-H. R. Chang, H. Swetha Koppula, and O. Tuzel, "Text is all you need: Personalizing asr models using controllable speech synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023, pp. 1–5.
- [16] T.-Y. Hu, M. Armandpour, A. Shrivastava, J.-H. R. Chang, H. Koppula, and O. Tuzel, "Synt++: Utilizing imperfect synthetic data to improve speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 7682–7686.
- [17] M. Le, A. Vyas, B. Shi *et al.*, "Voicebox: Text-guided multilingual universal speech generation at scale," in *Advances in Neural Information Processing Systems*, vol. 36, 2023, pp. 14 005–14 034.
- [18] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015, pp. 5206–5210.
- [19] K. Ito and L. Johnson, "The lj speech dataset," 2017.
- [20] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, "Hi-Fi Multi-Speaker English TTS Dataset," in *Interspeech*, 2021, pp. 2776–2780.
- [21] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," *CoRR*, vol. abs/1904.02882, 2019.
- [22] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *Oriental COCOSA 2017*, 2017, p. Submitted.
- [23] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "AISHELL-3: A multi-speaker mandarin TTS corpus and the baselines," *CoRR*, vol. abs/2010.11567, 2020.
- [24] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proc. 12th Conference on Language Resources and Evaluation*, 2020, pp. 4211–4215.
- [25] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, "Fleurs: Few-shot learning evaluation of universal representations of speech," 2022.
- [26] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, "SLURP: A spoken language understanding resource package," in *Proc. 2020 Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 7252–7262.
- [27] C. Veaux, J. Yamagishi, and K. MacDonald, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2019.
- [28] J. Park, S. Jin, J. Park *et al.*, "Conformer-based on-device streaming speech recognition with kd compression and two-pass architecture," in *IEEE Spoken Language Technology Workshop*, 2023, pp. 92–99.
- [29] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Transactions on Information and Systems*, vol. E99.D, no. 7, pp. 1877–1884, 2016.
- [30] J.-M. Valin and J. Skoglund, "A real-time wideband neural vocoder at 1.6 kb/s using LPCNet," in *Interspeech*, 2019.
- [31] Y. Wang, R. J. Skerry-Ryan, D. Stanton *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Interspeech*, 2017.
- [32] N. Ellinas, G. Vamvoukakis, K. Markopoulos *et al.*, "High quality streaming speech synthesis with low, sentence-length-independent latency," in *Interspeech*, 2020, pp. 2022–2026.

- [33] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," 2022.
- [34] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," 2019.
- [35] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proc. 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7871–7880.
- [36] J. Tiedemann and S. Thottingal, "OPUS-MT - Building open translation services for the World," in *Proc. 22nd Annual Conferenec of the European Association for Machine Translation*, 2020.
- [37] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022.
- [38] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," 2022.
- [39] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 17 022–17 033.
- [40] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019.
- [41] Y. Jia, Y. Zhang, R. Weiss *et al.*, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," in *Advances in Neural Information Processing Systems*, vol. 31, 2018.

Comparing Lazy Constraint Selection Strategies in Train Routing with Moving Block Control

Stefan Engels*  and Robert Wille*[†] 

*Chair for Design Automation, Technical University of Munich (TUM), 80333 Munich, Germany

[†]Software Competence Center Hagenberg GmbH (SCCH), 4232 Hagenberg, Austria

Email: {stefan.engels, robert.wille}@tum.de

Abstract—Railroad transportation plays a vital role in the future of sustainable mobility. Besides building new infrastructure, capacity can be improved by modern train control systems, e.g., based on moving blocks. At the same time, there is only limited work on how to optimally route trains using the potential gained by these systems. Recently, an initial approach for train routing with moving block control has been proposed to address this demand. However, detailed evaluations on so-called *lazy constraints* are missing, and no publicly available implementation exists. In this work, we close this gap by providing an extended approach as well as a flexible open-source implementation that can use different solving strategies. Using that, we experimentally evaluate what choices should be made when implementing a lazy constraint approach. The corresponding implementation and benchmarks are publicly available as part of the *Munich Train Control Toolkit* (MTCT) at <https://github.com/cda-tum/mtct>.

I. INTRODUCTION

SUSTAINABLE transportation systems are becoming increasingly important. Because of this, the demand for railway transportation is constantly increasing. Since building new tracks to increase capacity is resource- and time-consuming, train control systems should also be utilized to increase capacity.

Because trains cannot operate on sight due to their long braking distances, such systems are used to prevent collisions. Most notable systems are the *European Train Control System* (ETCS), the *Chinese Train Control System* (CTCT), or the *Positive Train Control* (PTC) [1] as well as *Communication Based Train Control* (CBTC) for metro systems [2]. While these systems differ in detail, the main concepts are very similar. New specifications allow trains to follow each other more closely on existing infrastructure and at the same level of safety. In the ideal case, trains can operate under a so-called moving block control, which provides enormous potential for increased capacity.

However, the most efficient specification does not help without methods to optimize train movements that use this potential. Respective optimization tasks using classical (i.e., “old”) specifications are well studied [3]. At the same time, there is only a little work on routing under moving block control [4], [5], none of which is available open-source.

Since the number of constraints preventing collisions is enormous and, at the same time, many of them are not needed to describe an optimal solution, a lazy approach is used. First,

the problem is optimized without these conditions. During the solving process, violated constraints are iteratively added until a feasible (hence, optimal) solution is obtained. There are different strategies on which (lazy) constraints to add in each iteration. However, to the best of our knowledge, they have not previously been compared, and it is hard to do corresponding evaluations ourselves due to the lack of available implementations.

In this work, we aim to improve upon the aforementioned. The resulting source code is included in the open-source *Munich Train Control Toolkit* (MTCT) available on GitHub at <https://github.com/cda-tum/mtct>. The solving strategy and other parameters can be chosen flexibly. This allows for experimental evaluations, in which we analyze what strategy should be implemented using a lazy approach in train routing under moving block. Additionally, the proposed model extends previous solutions to allow more general timetabling requests and can model train separation more precisely, especially in scenarios close to stations where a train might occupy multiple (short) train segments simultaneously.

The remainder of this work is structured as follows: Sec. II reviews the relevant principles of train control systems, Sec. III describes the considered routing task and summarizes previous work as well as our contribution, and Sec. IV and V present the proposed approach(es). Finally, Sec. VI contains an experimental evaluation, and Sec. VII concludes this paper.

II. TRAIN CONTROL PRINCIPLES

Classically, a railway network is divided into fixed blocks. Using *Trackside Train Detection* (TTD) hardware, e.g., *Axle Counters* (AC), it is determined whether a particular block is occupied or not. Because of this, the resulting blocks are also called TTD sections. A following train can only enter a block once it is fully cleared by the previous train.

Example II.1. Consider two trains following each other on a single track as depicted in Fig. 1a. Train tr_2 can only move until the end of TTD2. It cannot enter TTD3 because it is still occupied and, hence, might have to slow down in order to be able to come to a full stop before entering the occupied block section.

Modern control systems allow for more efficient headways. A train equipped with *Train Integrity Monitoring* (TIM) can




Fig. 1: Schematic drawings of different signaling principles

report its safe position to the control system. Hence, no hardware is needed to safely separate trains. Then, the network no longer has to be separated into fixed blocks. In the best case, trains can follow at absolute braking distance. Hence, shorter headway times are possible. This so-called *Moving Block* signaling has, e.g., been specified as an extension within ETCS Level 2 [6], which is also formerly known as Level 3 [7].

Example II.2. *In contrast to Ex. II.1, consider a moving block control implemented in Fig. 1b. Because trains operate at the ideal absolute braking distance, tr_2 can move up to the actual end of tr_1 (minus a little buffer). In particular, it can already enter what has been TTD3 previously. Hence, trains can follow each other more closely.*

III. PROBLEM DESCRIPTION AND CONTRIBUTIONS

In this work, we focus on moving block control systems. This section briefly provides the problem description of a corresponding routing task, reviews the current state of the art, and motivates our work.

A. Train Routing under Moving Block Control

Train routing is the task of determining when and where trains are driving on the respective network, given timetabling constraints. This includes the choice of specific tracks and corresponding timings. More formally, it is defined using the following notation:

\mathcal{T} : A set of trains and its relevant properties (e.g., length, maximal speed, acceleration, braking curves).

\mathcal{N} : A railway network including vertices V and (directed) edges E as described in [8].

\mathcal{S} : A set of stations, where each station $\mathcal{S} \ni S \subseteq E$ is a subset of edges of the network \mathcal{N} .

$\mathcal{D}^{(tr)}$: A set of demands for every train $tr \in \mathcal{T}$ consisting of

- a weight $w^{(tr)} \geq 0$ of importance,
- an entry node $v_{in}^{(tr)} \in V$ together with a desired entry interval $[\underline{t}_{in}^{(tr)}, \bar{t}_{in}^{(tr)}]$,
- an exit node $v_{out}^{(tr)} \in V$ together with a desired exit interval $[\underline{t}_{out}^{(tr)}, \bar{t}_{out}^{(tr)}]$, as well as,
- a set of stations $S_i^{(tr)} \in \mathcal{S}$ together with
 - * an interval $[\underline{\alpha}_i^{(tr)}, \bar{\alpha}_i^{(tr)}]$ in which the train should arrive at the station,
 - * an interval $[\underline{\delta}_i^{(tr)}, \bar{\delta}_i^{(tr)}]$ in which the train should depart from the station, and
 - * a minimal stopping time $\Delta t_i^{(tr)} \geq 0$.

Having this notation, the goal is to determine an optimal routing. In this setting, optimality is defined as minimizing the (weighted) exit times such that all schedule demands are obeyed and the constraints by a moving block control system are satisfied.

B. State of the Art and Contributions

Train routing and related timetabling tasks under classical train control have long been considered and are well-studied [3]. On modern control systems using so-called hybrid train detection, routing is considered in algorithms to design optimal (virtual) section layouts by using SAT [9], A* [10], or *Mixed Integer Linear Programming* (MILP) [11], [12]. While the arising questions are similar, these solutions do not utilize the full potential of moving block.

To the best of our knowledge, [4] is the first approach that considers optimal routing of trains specifically under moving block control. They describe a MILP formulation to solve a routing problem similar to the one considered in this paper. Say s describes the position and t the time; one could say that their formulation models the function $t(s)$ at discrete positions given by vertices of the network. Since trains cannot pass each other on a given edge, this seems to be a reasonable simplification while still being able to model at a decent level of accuracy.

However, the number of constraints to ensure that trains keep enough distance and do not crash into each other is rather big. At the same time, most of these are unnecessary because trains operating at different network parts will not collide even without explicit constraints. Because of this, one can first optimize without them. If this yields a collision-free solution, the problem is solved. Otherwise, constraints preventing collisions from arising need to be added during the solving process as so-called *lazy constraints*. By doing this, the same optimal solution is obtained; however, only a small number of the original constraints is considered. This can be beneficial, especially for large models, as discussed in their follow-up work [5].

At the same time, this previous approach comes with some downsides:

- Both trains and stations are single points without length. The authors claim this is not a problem because the length can be integrated as a buffer in the headway. However, especially in station environments, this might not be feasible. For example, some stations separate a platform into sections. A long train might occupy all of a platform, whereas two shorter trains can stop simultaneously (in different sections of the platform). Those scenarios cannot be modeled using the previous approach.
- There are different strategies to select which (lazy) constraints to add. This constitutes a trade-off: adding only a few lazy constraints in each iteration is quickly possible. However, many iterations might be needed until a collision-free solution is reported. On the other hand, simultaneously adding many lazy constraints increases the time spent in every iteration but, at the same time,

likely reduces the number of iterations needed. In [5], no evaluation of selection strategies for lazy constraints is provided.

- The implementation of the solution is not publicly available. This prevents us from doing corresponding evaluations and restricts the proposed solution's applicability.

Overall, this motivates an alternative MILP formulation, which properly considers train separation even on shorter edges by considering the respective train lengths as well as incorporating more flexible timetable requests. At the same time, we aim to shed light on which strategy for lazy constraint selection might be best by conducting a corresponding evaluation. Finally, we provide a flexible open-source implementation at <https://github.com/cda-tum/mtct>, thus allowing the community to access such methods.

IV. MILP MODEL

This section presents the MILP model motivated in Sec. III-B. For reasons of comprehensibility, we limit ourselves to the relevant variables and constraints. The interested reader can find the complete model in the open-source implementation available at <https://github.com/cda-tum/mtct>.

A. Symbolic Formulation

To model the approach, we need variables describing each train's routes and relevant timings. As discussed in Sec. III-B, we follow the basic strategy by [4] with slight extensions to incorporate the actual train lengths by tracking each train's rear point. Hence, we include the following variables:

- $x_e^{(tr)} \in \{0, 1\}$ denotes whether a certain edge $e \in E$ is used by train $tr \in \mathcal{T}$.
- $\bar{a}_v^{(tr)} \in [0, \bar{t}_{out}^{(tr)}]$ is the time at which the front of train $tr \in \mathcal{T}$ arrives at $v \in V$.
- $\bar{d}_v^{(tr)} \in [0, \bar{t}_{out}^{(tr)}]$ is the time at which the front of train $tr \in \mathcal{T}$ departs from $v \in V$.
- $\underline{d}_v^{(tr)} \in [0, \bar{t}_{out}^{(tr)}]$ is the time at which the rear of train $tr \in \mathcal{T}$ departs from $v \in V$, hence, tr has entirely left the previous edge.

The speed is included by extending the vertices accordingly. Let $\mathcal{P}_v^{(tr)} \subseteq [0, v_{max}^{(tr)}]$ be a finite set of discretized velocities (paces) a train $tr \in \mathcal{T}$ might have at $v \in V^1$. The extended graph has a (directed) edge $\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \in E^{(tr)}$ connecting (u, p_1) and (v, p_2) for $e = (u, v) \in E$, $p_1 \in \mathcal{P}_u^{(tr)}$, and $p_2 \in \mathcal{P}_v^{(tr)}$ if, and only if, it is possible for train $tr \in \mathcal{T}$ to accelerate/decelerate from p_1 to p_2 while traveling on e . For every such extended edge $\epsilon_{e,p_1 \rightarrow p_2}^{(tr)}$, the variable

- $y_{e,p_1 \rightarrow p_2}^{(tr)} \in \{0, 1\}$ denotes whether train $tr \in \mathcal{T}$ uses the corresponding extended edge.

Example IV.1. Consider again the setting of Ex. II.2, where two trains follow each other. For presentation purposes, we choose the segment to consist of three vertices. In Fig. 2, the

¹For this work, we have used a uniform discretization of 10km/h as these values can be displayed by a standard speed indicator [13, Signal Zs 3].

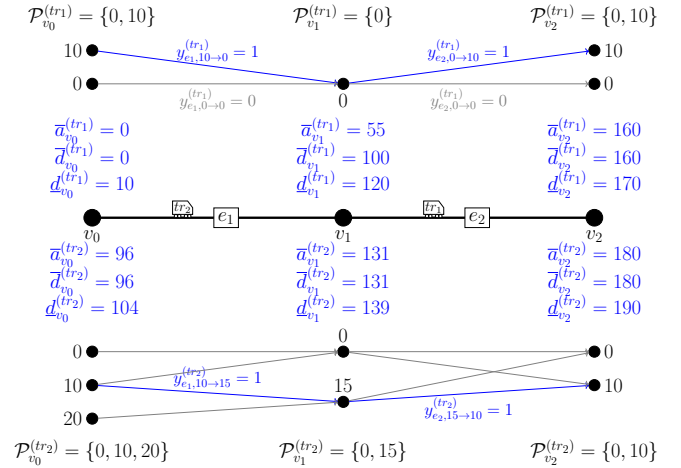


Fig. 2: Example Setting for Symbolic Formulation

extended graph for tr_1 is drawn above and the one for tr_2 below the track. Values of relevant variables are written at the fitting places. Furthermore, note that $x_{e_1}^{(tr_1)} = x_{e_2}^{(tr_1)} = x_{e_1}^{(tr_2)} = x_{e_2}^{(tr_2)} = 1$.

B. Constraints

Of course, the above variables must be constrained to represent valid train movements. In the following, we present the most important constraints. We leave out some technical constraints needed to connect variables according to their desired purpose for ease of understanding. Furthermore, we might formulate some constraints in its logical, rather than linear, form. In either case, adding and reformulating the missing constraints using big-M and possibly some helper variables is straightforward. They are irrelevant to presenting this work's crucial concepts and ideas.

1) *Valid Train Movements:* Each train $tr \in \mathcal{T}$ has to travel on a valid path from its entry to exit. Hence,

$$\sum_{e \in \delta^+(v_{in}^{(tr)})} x_e^{(tr)} = \sum_{e \in \delta^-(v_{out}^{(tr)})} x_e^{(tr)} = 1 \quad (1)$$

where $\delta^+(\cdot)$ and $\delta^-(\cdot)$ denote outgoing and incoming edges respectively. Moreover, a flow-conserving constraint (which also considers the respective velocity) has to be fulfilled at every other vertex. Thus, for every $v \in V - \{v_{in}^{(tr)}, v_{out}^{(tr)}\}$ and $p \in \mathcal{P}_v^{(tr)}$ we have

$$\sum_{\epsilon \in \delta^-(v,p)} y_\epsilon^{(tr)} = \sum_{\epsilon \in \delta^+(v,p)} y_\epsilon^{(tr)} \quad (2)$$

within the velocity extended graph. To prevent cycles, the in- and out-degrees are furthermore bound by one, respectively.

2) *Travel Times:* Denote by $\underline{t}^{(tr)}(\epsilon_{e,p_1 \rightarrow p_2}^{(tr)}) \in \mathbb{R}_{\geq 0}$ the minimal time it takes train $tr \in \mathcal{T}$ to traverse $\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \in E^{(tr)}$ of the velocity-extended graph. For details on how these may be calculated, we refer to [4, Fig. 3], but it suffices

to consider them as an arbitrary oracle. Analogously, let $\bar{\tau}^{(tr)} \left(\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \right) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ denote the maximal time. In this case, it is noted that a train might be allowed to stop on some of the edges, in which case ∞ is possible. Again, we refer to [4, Fig. 4].

Hence, assuming $e = (u, v)$, we have

$$\bar{a}_v^{(tr)} \leq \bar{a}_u^{(v)} + \underline{\tau}^{(tr)} \left(\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \right) + M \cdot \left(1 - y_{e,p_1 \rightarrow p_2}^{(tr)} \right) \quad (3)$$

$$\bar{a}_v^{(tr)} \geq \bar{a}_u^{(v)} + \bar{\tau}^{(tr)} \left(\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \right) - M \cdot \left(1 - y_{e,p_1 \rightarrow p_2}^{(tr)} \right) \quad (4)$$

where $M \geq 0$ is large enough (e.g., $M = \bar{t}_{out}^{(tr)2}$) to ensure that the constraint is only activated if the respective edge is used.

Finally,

$$\bar{d}_v^{(tr)} \geq \bar{a}_v^{(tr)} \quad \forall v \in V, tr \in \mathcal{T} \quad (5)$$

and a train can only stop at a vertex (i.e., “ \neq ” in Eq. (5)) if it has velocity 0, i.e.,

$$\bar{d}_v^{(tr)} \leq \bar{a}_v^{(tr)} + M \cdot \sum_{\epsilon \in \delta(v,0)} y_{\epsilon}^{(tr)}. \quad (6)$$

3) *Track Release*: In contrast to [4], we do not model a train as a single point. This allows for more accurate train separation in the model. For this, we need to relate the end of a train to its front. Let $\mathcal{R} = \{e_1, \dots, e_k\} \subseteq E$ be a route starting in v , such that $\sum_{i=1}^{k-1} l(e_i) < l(tr) \leq \sum_{i=1}^k l(e_i)$ ³. Similarly to above, let $\underline{\tau}_{\lambda \rightarrow \mu}^{(tr)}(\epsilon)$ and $\bar{\tau}_{\lambda \rightarrow \mu}^{(tr)}(\epsilon)$ be the minimal and maximal travel time from point λ to μ on the velocity extended edge $\epsilon \in E^{(tr)}$, where $0 \leq \lambda \leq \mu \leq l(\epsilon)$. Then, the following bounds have to hold

$$x_e^{(tr)} = 1 \forall e \in \mathcal{R} \Rightarrow \underline{d}_{u_1}^{(tr)} \geq \bar{a}_{u_k}^{(tr)} + \sum_{\epsilon \in \mathcal{E}_k} y_{\epsilon}^{(tr)} \cdot \underline{\tau}_{0 \rightarrow s}^{(tr)}(\epsilon) \quad (7)$$

$$x_e^{(tr)} = 1 \forall e \in \mathcal{R} \Rightarrow \underline{d}_{u_1}^{(tr)} \geq \bar{a}_{v_k}^{(tr)} - \sum_{\epsilon \in \mathcal{E}_k} y_{\epsilon}^{(tr)} \cdot \bar{\tau}_{s \rightarrow l_k}^{(tr)}(\epsilon) \quad (8)$$

assuming $e_i = (u_i, v_i)$, $s := l(tr) - \sum_{i=1}^{k-1} l(e_i)$, $l_k := l(e_k)$, and \mathcal{E}_k being the set of all edges connecting u_k to v_k in the velocity extended graph.

While we chose to write the logical form in Eq. (7) and (8) for better readability, they can easily be reformulated into linear constraints using big-M. We do not add upper bounds because the objective of small headways pushes the variables down wherever needed.

4) *Headway*: Reference [4] models train headways on single edges, which is precise if edges are rather long. However, the braking distance considered might range multiple edges, particularly close to stations. We use and proceed similarly to Sec. IV-B3 to model this more precisely. However, the length of the train is replaced by its braking distance.

On each edge $e \in E$, we introduce binary variables $o_e^{tr_1 \succ tr_2} \in \{0,1\}$, which is 1 if, and only if, $tr_1 \in \mathcal{T}$

²Note that Eq. (4) is only added if $\bar{\tau}^{(tr)} \left(\epsilon_{e,p_1 \rightarrow p_2}^{(tr)} \right) \leq \bar{t}_{out}^{(tr)}$ because otherwise bounding by the maximal travel time has no effect.

³In general we denote by $l(\cdot)$ the length of an object

follows $tr_2 \in \mathcal{T}$ on edge e . The respective headway constraints relating \bar{a} of the following and \underline{d} of the preceding train are then analog to Eq. (7) and (8); however, with the additional conditions that the following train has a specific velocity and the respective ordering variable is one.

Similarly, one can proceed with trains traveling in opposite directions. Then, however, the respective track segments behave like a TTD section, and a train’s moving authority can only enter a track segment once the opposing train has entirely left it.

5) *Timetable*: Of course, also the timetable demands $\mathcal{D}^{(tr)}$ have to be satisfied. Reference [4] can bind the respective timing variables directly since the exact stopping points are predefined. While, in our case, this is true for the entry and exit nodes, each stop at station $S_i^{(tr)} \in \mathcal{S}$ could be at a particular set of vertices, say $V_{S_i}^{(tr)}$. For every such $v \in V_{S_i}^{(tr)}$, we add a respective binary variable $stop_{i,v}^{(tr)} \in \{0,1\}$. Then,

$$stop_{i,v}^{(tr)} = 1 \Rightarrow \bar{a}_v^{(tr)} \in [\underline{\alpha}_i^{(tr)}, \bar{\alpha}_i^{(tr)}], \quad (9)$$

$$stop_{i,v}^{(tr)} = 1 \Rightarrow \bar{d}_v^{(tr)} \in [\underline{\delta}_i^{(tr)}, \bar{\delta}_i^{(tr)}], \text{ and} \quad (10)$$

$$stop_{i,v}^{(tr)} = 1 \Rightarrow \bar{d}_v^{(tr)} - \bar{a}_v^{(tr)} \geq \Delta t_i^{(tr)}. \quad (11)$$

Again, these logical constraints can easily be reformulated into linear constraints using big-M.

C. Objective

Finally, the goal is to enable every train to leave the network as early as possible. If a train leaves after its predefined earliest departure time, it is caused by the routing choice, not the respective request. We minimize this difference according to the given weights, which we normalize to one. Thus, the objective is given by

$$\min \frac{1}{\sum_{tr \in \mathcal{T}} w^{(tr)}} \cdot \sum_{tr \in \mathcal{T}} w^{(tr)} \cdot \left(\underline{d}_{v_{out}}^{(tr)} - \underline{t}_{out}^{(tr)} \right). \quad (12)$$

V. LAZY HEADWAY CONSTRAINTS

Note that there are many headway constraints of the form described in Sec. IV-B4, more precisely of order $\mathcal{O} \left(|\mathcal{T}| \cdot \sum_{tr \in \mathcal{T}} \sum_{v \in V} |\mathcal{P}_v^{(tr)}| \right) = \mathcal{O} \left(|\mathcal{T}|^2 \cdot |V| \cdot |\overline{\mathcal{P}}| \right)$, where $|\overline{\mathcal{P}}|$ denotes the average number of velocity extensions. For instances with many trains on more extensive networks, the time to explicitly add all these constraints to a model is substantial. However, most of these constraints are not explicitly needed because they describe a scenario far from optimal. This motivates a lazy approach.

For this, we optimize using all except the headway constraints. The obtained solution could violate some of the requirements. If so, one has to add a set of violated constraints to the model and reoptimize, whereby the solver can use information from the previous iteration to warm start. This procedure is continued until the solution is feasible and, hence, optimal.

However, the question arises of which constraints to add in each iteration. A given conjectured solution determines each




Fig. 3: Runtime of Different Lazy Constraint Strategies

train’s route and velocity profile uniquely. Hence, at most $\mathcal{O}(|\mathcal{T}| \cdot \sum_{tr \in \mathcal{T}} |\mathcal{X}^{(tr)}|) = \mathcal{O}(|\mathcal{T}|^2 \cdot \overline{|\mathcal{X}|})$ conditions need to be checked, where $|\mathcal{X}^{(tr)}|$ denotes the number of edges used by a respective train and, again, $\overline{|\mathcal{X}|}$ the respective average. In fact, the number is even smaller because only trains whose routes intersect need to be compared.

Note that the headway constraints are transitive in the sense that if tr_1 follows tr_2 at a safe distance at a given point, then it also safely follows all trains that might have passed before tr_2 . In particular, it suffices to check (and possibly add) only headways regarding the immediate preceding train. This reduces the number of checked constraints to $\mathcal{O}(|\mathcal{T}| \cdot \overline{|\mathcal{X}|})$.

Finally, one could either add all checked constraints (to give as much information as possible to the solver) or only such constraints that are violated (to not overload the solver with unnecessary information). In the extreme case, one could even only add one violated constraint and stop checking further constraints immediately. In the first case, one is likely to only need a few iterations. Conversely, a first violated constraint might be found very quickly; however, more iterations are needed in the end. The best strategy might depend on the specific problem and is worth evaluating.

VI. IMPLEMENTATION AND EVALUATION

The approach presented above has been implemented, made publicly available as open-source, and used to evaluate lazy constraint selection strategies. It is included in the *Munich Train Control Toolkit* available at <https://github.com/cda-tum/mtct> and under active development. In this section, we describe both the resulting implementation as well as the evaluations and results obtained.

A. Implementation

We implemented the model described in Sec. IV using the C++ API of Gurobi [14]. The resulting tool allows the user to choose between different strategies for lazy constraints to be added in each iteration by controlling various parameters.

Lazy constraints are implemented using the (custom) callback framework provided by Gurobi.

In the current version, the tool allows to, e.g., compare the following selection strategies:

- “*Full Model*:” The entire model is explicitly constructed in advance and passed to the solver. No callback is used.
- “*All Checked Lazy Constraints*:” In case of infeasibility, all $\mathcal{O}(|\mathcal{T}|^2 \cdot \overline{|\mathcal{X}|})$ constraints corresponding to overlapping routes are added in each iteration regardless if they are violated or not.
- “*All Adjacent Train Constraints*:” Similarly, all checked constraints are added in case of infeasibility. However, only $\mathcal{O}(|\mathcal{T}| \cdot \overline{|\mathcal{X}|})$ conditions corresponding to adjacent trains directly following each other are considered.
- “*Adjacent Violated Constraints*:” Again, only constraints corresponding to adjacent trains are considered. This time, however, only violated constraints are passed to the solver in each iteration. Conditions already fulfilled are ignored but might be added to a later callback.
- “*Only First Violation*:” As soon as one violated constraint is found, only this one is added, and the callback is immediately aborted without checking the remaining conditions.

B. Evaluation

We tested these different strategies on an Intel(R) Xeon(R) W-1370P system using a 3.60GHz CPU (8 cores) and 128GB RAM running Ubuntu 20.04 and Gurobi version 11.0.2. As benchmarks, we use the railway networks and schedules from [11, Appendix A]. Additionally, we create random timetables of up to 50 trains on two of the networks, including the Munich S-Bahn Stammstrecke. Since optimizing up to the millisecond is unreasonable, we stop at a proven optimality gap of 10 seconds.

The results are provided in Fig. 3. On the x-axis, we plot the runtimes in seconds. Note that we chose a logarithmic scale for better readability. The y-axis provides the fraction of samples that were solved in the given time or faster. The

lines are monotonously increasing by design. Generally, if a line is over/left of another line, the corresponding algorithm performs faster/better.

We present two plots to avoid distorting the analysis because of infeasible instances. On the left, we included instances known to be feasible; on the right, we included instances proven to be infeasible. In the latter case, the time plotted corresponds to the time it took the proposed approach to prove infeasibility.

Clearly, the numbers confirm that a lazy approach is beneficial: one should not explicitly specify the entire model in advance. The only exception to this is that only adding one constraint at a time performs even worse, which becomes especially clear when considering infeasible examples.

Among the other strategies, no one clearly outperforms the others. At the same time, there seems to be a slight advantage of only considering adjacent trains directly following each other instead of all possible pairs of trains. However, it is questionable if this effect is significant.

Overall, it seems reasonable to only add violated constraints corresponding to adjacent trains. However, other strategies might also be beneficial depending on the context in which the algorithm is used since the observed benefit is only minor.

Having the proposed approach available as open-source will allow adding and evaluating further strategies easily.

VII. CONCLUSIONS

In this work, we considered train routing within a moving block environment. We introduced a MILP formulation that can more accurately (than existing solution methods) model train separation on layouts with short track segments by incorporating the actual train length. Moreover, we discussed how a lazy constraint approach can be implemented using different strategies in each callback. Various such strategies have been implemented open-source and are available at <https://github.com/cda-tum/mtct>. The user can control the parameters affecting the solving process.

An experimental evaluation confirms that the solution process benefits from the lazy approach as long as multiple constraints are added simultaneously. On the other hand, there seems to be no significant difference between some of the tested strategies. At the same time, the open source implementation allows for the use of different strategies depending on the instance, and it is not necessary to decide on the one and only best approach in this setting.

Previous work focuses on the optimal design of other modern train control systems relying on so-called hybrid train detection. These systems combine the efficiency of moving block with the practicability of classical train control [15]. Design automation methods in this context must both route trains and place so-called virtual subsections. However, both of these tasks alone are already hard, and optimization methods in this context can highly benefit if routing is considered separately [11]. While the details are out of scope for this paper, it is reasonable to believe that optimal routes under moving block are good choices in this case. In particular,

we aim to include this work (and possible future work on routing under moving block control) as a first step within an optimization pipeline for automated planning of train control systems with hybrid train detection. Again, future work will also be made available open-source as part of the Munich Train Control Toolkit mentioned above.

REFERENCES

- [1] J. Pahl, *Railway Signalling Principles: Edition 2.0*, 2021. [Online]. Available: <http://dx.doi.org/10.24355/dbbs.084-202110181429-0>
- [2] L. Schnieder, *Communications-Based Train Control (CBTC)*. Springer Berlin Heidelberg, 2021. [Online]. Available: <http://dx.doi.org/10.1007/978-3-662-62876-8>
- [3] R. Borndörfer, T. Klug, L. Lamorgese, C. Mannino, M. Reuther, and T. Schlechte, Eds., *Handbook of Optimization in the Railway Industry*, 2018. [Online]. Available: <http://dx.doi.org/10.1007/978-3-319-72153-8>
- [4] T. Schlechte, R. Borndörfer, J. Denißen, S. Heller, T. Klug, M. Küpper, N. Lindner, M. Reuther, A. Söhlke, and W. Steadman, "Timetable optimization for a moving block system." *Journal of Rail Transport Planning & Management*, vol. 22, 2022. [Online]. Available: <http://dx.doi.org/10.1016/j.jrtpm.2022.100315>
- [5] T. Klug, M. Reuther, and T. Schlechte, "Does laziness pay off? - a lazy-constraint approach to timetabling," in *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, 2022. [Online]. Available: <http://dx.doi.org/10.4230/OASICS.ATMOS.2022.11>
- [6] Siemens, Bombardier, Mermec, Network Rail, and Thales, "Deliverable D4.2 moving block enhancements," in *X2Rail-5 Completion of activities for Adaptable Communication, Moving Block, Fail Safe Train Localisation (including satellite), Zero on site Testing, Formal Methods and Cyber Security*. Shift2Rail, 2023. [Online]. Available: https://projects.shift2rail.org/s2r_ip2_n.aspx?p=X2RAIL-5
- [7] Siemens, Hitachi Rail STS, Bombardier, Thales, Network Rail, Alstom, CAF, Trafikverket, AZD, Mermec, Deutsche Bahn, SNCF, and ERTMS Users Group, "Deliverable D5.1 moving block system specification," in *X2Rail-1 Start-up activities for Advanced Signalling and Automation Systems*. Shift2Rail, 2019. [Online]. Available: https://projects.shift2rail.org/s2r_ip2_n.aspx?p=X2RAIL-1
- [8] S. Engels, T. Peham, J. Przigoda, N. Przigoda, and R. Wille, "Design tasks and their complexity for the European Train Control System with hybrid train detection," 2024. [Online]. Available: <http://dx.doi.org/10.48550/arXiv.2308.02572>
- [9] R. Wille, T. Peham, J. Przigoda, and N. Przigoda, "Towards automatic design and verification for Level 3 of the European Train Control System," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021. [Online]. Available: <http://dx.doi.org/10.23919/date51398.2021.9473935>
- [10] T. Peham, J. Przigoda, N. Przigoda, and R. Wille, "Optimal railway routing using virtual subsections," in *Reliability, Safety, and Security of Railway Systems. Modelling, Analysis, Verification, and Certification*, 2022. [Online]. Available: http://dx.doi.org/10.1007/978-3-031-05814-1_5
- [11] S. Engels, T. Peham, and R. Wille, "A symbolic design method for ETCS Hybrid Level 3 at different degrees of accuracy," in *23rd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, 2023. [Online]. Available: <http://dx.doi.org/10.4230/OASICS.ATMOS.2023.6>
- [12] S. Engels and R. Wille, "Late breaking results: Iterative design automation for train control with hybrid train detection," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024. [Online]. Available: <http://dx.doi.org/10.23919/DATE58400.2024.10546590>
- [13] DB InfraGO AG, "Richtlinie 301: Signalbuch," 2024. [Online]. Available: https://www.dbinfrago.com/web/schienennetz/netzzugang-und-regulierung/regelwerke/betrieblich-technisch_regelwerke
- [14] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [15] M. Bartholomeus, L. Arenas, R. Treydel, F. Hausmann, N. Geduhn, and A. Bossy, "ERTMS Hybrid Level 3," *SIGNAL + DRAHT (110) I+2*, 2018. [Online]. Available: https://www.eurailpress.de/fileadmin/user_upload/SD_1_2-2018_Bartholomeaus_ua.pdf

Dynamic Threat Intelligence for Improvement of Resilience of Critical Infrastructure During Pandemics

Pablo de Juan Fidalgo
0000-0003-1347-6000
Eviden Spain, Albarracin 25,
28037 Madrid, Spain, Email:
pablo.dejuan@eviden.com

Aljosa Pasic
0000-0003-0150-5732
Eviden Spain, Albarracin 25,
28037 Madrid, Spain
Email: aljosa.pasic@eviden.com

Susana Gonzalez Zarzosa
0000-0002-3402-2385
Eviden Spain, Albarracin 25,
28037 Madrid, Spain, Email:
susana.gzarsosa@eviden.com

Abstract— The COVID-19 pandemic is an example of a temporary situation when critical infrastructure (CI) operators had to operate with continuously changing conditions. The role of cyber infrastructure during pandemics, for example for the remote work or access to critical systems, has also changed. This resulted in frequent re-evaluation of risks and adaptations of security policies or mitigation measures. Use and sharing of cyber threat intelligence (CTI) proved to be valuable to stay up to date, but challenges related to trust and confidence emerged. We designed and developed dynamic CTI to be used by CI operators for risk reassessment and improvement of resilience. Several enhancements will be validated in the forthcoming pilots in SUNRISE project.

Index Terms—Cybersecurity, cyber resilience, threat intelligence, critical infrastructure, pandemics.

I. INTRODUCTION

RESILIENCE is a concept that originally comes from disaster and crisis management and is closely related to efforts to analyse and manage preparedness, but also resistance and recovery from adverse events. It refers to the ability of a system, community, or society to resist, absorb, accommodate to, and recover from the effects of a hazard or incident. It is also characterized by “changing conditions” and “deliberate attacks”, especially if we move to cybersecurity resilience definitions [1, 2]. When it comes to general characteristics of resilience, described in [3], these include uncertainty and dynamism in different circumstances, as well as complexity and difficulty to measure. These are especially challenging for the resilience of complex and interconnected systems such as Critical Infrastructures (CIs). In this context, we can mention many EU efforts, starting from the European program for critical infrastructure protection (Directive 2008/114 [4]), which establishes a procedure for identifying and designating European CI, to a more recent Directive on the Resilience of Critical Entities (CER Directive, PE-CONS 51/22 [5]) that entered into force on 16 January 2023. There are several slightly different definitions of CIs or Critical Entities (CEs) but for the

purpose of simplicity (and because Member States have until 17 October 2024 to adopt national legislation to transpose the CER Directive), we will keep term CIs. In contrast to the previous approach, with more focus on prevention and mitigation, CER directive also focuses on the response and the rapidity of recovery during and after the event.

During the COVID-19 pandemic, we have seen what the consequences are not only for the healthcare system, but equally as much for other interdependent CI entities. The functioning of CIs was highly unpredictable due to the multitude interdependencies, including interruptions in supply chain, effects on essential employees both physically and mentally, or changes of organizational priority and incident response strategy. To address these challenges, SUNRISE research and innovation project was co-funded by the European Commission and started in November 2022 with a duration of 36 months. Project activities are executed by a consortium of 42 partners from different countries and are considering research activities, as well as the integration and adaptation of relevant tools and practices related to CI resilience, based on the lessons learned from pandemics, such as interlinking of different types of risk inputs and indicators. The SUNRISE strategy consists of distinctive missions, one of them being awareness of the dynamic threat landscape related to and implied by pandemics. One of the key results of SUNRISE project to accomplish this mission is an integral cyber-physical resilience (CPR) tool composed of:

- AI-based anomaly detector (AD) that analyses log files from different CI systems and components.
- Enhanced cyber threat intelligence (CTI) platform and sharing service for (CI) operators.
- Semi-automated risk assessment engine, extended to include physical risk indicators, models that correspond to pandemic-specific strategies of CI operators, as well as temporary condition changes.
- Incident response and reporting management that takes inputs from other tools, including legacy tools, and considers incident reporting thresholds and

workflows that are compliant with EU and national legislation for CI.

There are several innovations in each part of CPR, which are also described in a previous paper about CPR architecture [6]. In this paper we will focus on concepts and design of enhancements of CTI platform, some of which are related to pandemics scenarios of adaptivity, collaboration and absenteeism, while others are applicable to CTI in any context and conditions. We start with the description of related work, before addressing conceptual design and implementation of enhancements. Since these are related to threat intelligence and model, we also explain links to the main risk and resilience concepts, such as calculation of probabilities. We briefly cover validation, which is still in progress and finish with conclusions.

II. RELATED WORK

ENISA launched already in 2010, a study on Measurement Frameworks and Metrics for Resilient Networks and Services. The methodology included both a survey, a desktop research and further consultation [7], as well as supportive taxonomy issued in 2011 [8]. In [9] authors deal with challenge of representing interconnections among system components across operational domains (physical, information, cognitive, or social) and present the cyber resilience matrix with four capacities (plan and prepare/absorb/recover/adapt) and four domains. None of these works addresses the use and sharing of threat intelligence and its impact on resilience.

In some of our previous works ([10],[11],[12]) several enhancements of threat intelligence platform, and enhancement of operational threat indicators have been described. The platform receives structured cyber threat information from multiple sources and performs the correlation with both static and dynamic data coming from the monitored infrastructure. This allows the evaluation of a threat score through heuristic-based analysis, used for enriching the information received from open-source intelligence (OSINT) and other sources. While these remain relevant and useful, in fast changing temporary conditions, such as during pandemics, there is a need for further adaptation to have dynamic adaptation of parameters used threat intelligence trust scoring or source confidence.

In addition, experience with COVID-19 showed that both closed community-based platforms, as well as social media platforms were used to share news, ideas, or opinions, which enabled data processing to serve more efficient and valuable pandemics surveillance. Data derived from related health trends can help to predict workforce availability or absenteeism, and in this way improve critical infrastructure (CI) risk assessment during pandemics. Collaboration between various public health stakeholders, for example, is exemplified in Epidemic Intelligence from Open Sources (EIOS) initiative [13], regional monitoring [14] or European Centre for disease prevention and control (ECDC) project that performs epidemic intelligence [15] for early detection, verification, assessment, and communication of health

threats. The use of crowdsourced OSINT as an alternative to commercial or community-based threat intelligence, is also explored and described in many papers, for example its effectiveness in malware detection which is described in [16], or crowdsourced support for discovery and verification of OSINT sources [17]. Confidence in sources and threat score proved to be one of the biggest challenges, especially due to the context and dynamics of temporary situations.

In our work we use open-source Malware Information Sharing Platform (MISP) platform [18] as a basis for our CTI tool. There was already work on COVID-19 MISP instance focusing on three areas of sharing: medical information, cyber threats related to COVID-19 and disinformation about COVID-19. In addition, MISP taxonomies [19] and tools for mapping and a comprehensive checklist of activities related to MISP implementation in the context of COVID-19 that cut across critical domains [20], have also been developed and offered free of charge. This information sharing community has a low barrier of entry, and everyone can contribute or use the data, which is one of the reasons why in SUNRISE we conceptualized and developed CTI tool with fine grained data sharing policy, as well as dynamic threat scoring and source confidence detection.

III. CPR DESIGN AND CTI TOOL ENHANCEMENTS

In this chapter, our focus of enhancements is threat intelligence scoring module, that was extended with source confidence evaluation for better management of intelligence sources, mapping of Indicators of Compromise (IoCs) with MITRE ATT&CK matrix techniques and sharing of contextual events and temporary conditions that are related and relevant for pandemics (e.g. absenteeism, workforce availability prediction, etc). However, since CTI is used as a module within CPR, we will also describe some of CPR concepts.

A. Conceptual Design

During temporary unforeseen circumstances like pandemics, the availability of qualified staff may fluctuate, potentially impacting the credibility of certain sources within the CTI ecosystem. This fluctuation exemplifies the necessity for adaptivity, where systems must dynamically adjust their trust assessments based on evolving contextual factors. The dynamic adjustment of score confidence allows CTI systems to recalibrate their trust scores in real-time, accounting for changes in source reliability and relevance. We argue that for the fast-changing cybersecurity landscape, and contextualization of the available information in changing circumstances, approach based on Observe, Orient, Decide, Act (OODA) loop is more appropriate than predominant Plan, Do, Check, Act (PDCA) approach.

In the practical application of this approach, we start from what is available among observation sources: types of data or assessments that already exist in CI operators that act as SUNRISE project users. In parallel we develop an idea about what is desirable for CPR tool in general and CTI

enhancement in particular, especially for the orientation and decision phase of OODA loop. We proceed with a theoretic model that links different observed or inferred data/events (from log files, network, external threat sources, surveys, manual inputs etc.) with two level orientation and decision assessments (Figure 1):

- strategic/tactical level, dealing with priorities, asset values etc
- operational level, with rule-based assessments and incident response

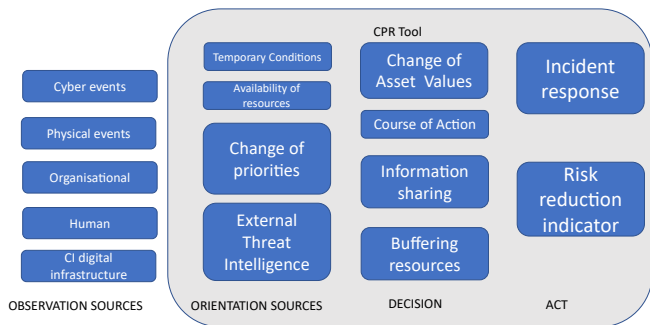


Fig 1. Mapping of CPR tool Conceptual Design to OODA loop

There are many possible enhancements in orientation phase, where external threat intelligence can be used. The effectiveness of security controls may decrease due to temporary changes in another CI (e.g. due to the disruption of supply chain or workforce availability). The threat environment or the technical environment (e.g., the introduction of new technologies in CI digital infrastructure such as pandemic-specific access control equipment) also change. In parallel, the relative priorities of cyber resiliency objectives may shift based on inputs from other stakeholders, current concerns, or available resources and funding.

B. Link of CTI to Risk Assessment and Threat Model

The risk assessment module takes into consideration several types of inputs, both static and real-time. These come from observation and orientation sources. For example, inputs related to the operational context model (business profile, sector specific inputs such as supply chain, human resources etc) are static and come from the relevant employees. Target asset model also uses static inputs about e.g. asset value, asset connectivity etc. Both types of inputs have separated user interfaces, but these can be changed through a new option in dashboard for “temporary conditions” which overwrites the default values. Application, which collects health-related data of citizens, will typically have higher asset value in temporary model, which overwrites its default value, and therefore increases impact and related risks. Similar holds for absenteeism prediction input that might increase vulnerabilities value in assets operated by absent employees, and therefore related risks would also increase.

Unlike static inputs, real data ingestion model does not depend on temporary conditions. This data is provided by the existing cybersecurity tools, such as intrusion detection system (IDS), security event and information management system (SIEM), physical access events and others.

Selection of threat models was done by using MITRE ATT&CK knowledge base of adversary tactics and techniques based on real-world observations. ATT&CK framework is also used to build detailed attack emulation scenarios and to verify that the events collected by agents and analytics can be used as risk indicators and mapped into risk models. In the case of a highly protected and controlled environment, such as in CI, tactic such as privilege escalation can be an objective with associated technique such as access token manipulation or theft. With escalated privileges, an adversary could, for example, program the mobile device to impersonate USB devices such as input devices (keyboard and mouse), storage devices, and/or networking devices. Other examples of physical events and human errors could include attempt to exfiltrate data or insert ransomware/malware over a USB connected to a physical device, or “shoulder surfing”, which in SUNRISE is detected by means of video processing and is fed into the risk assessment module. Combination with IoC received about phishing attacks in organisations that work as a supplier to CI operator, could, for example, reveal that attack on CI core system is in preparation.

Each threat is usually associated with several inputs, that could be further linked to specific questions for questionnaire or detection by some cybersecurity tools. For example, Phishing and Impersonation Attacks questions are: Are employees aware and trained against Phishing? Is multi-factor authentication (MFA) implemented for user login? Is the account locked or did you have too many login attempts? Is there an unusual user behaviour or pattern?

C. Threat Intelligence Engine Architecture

In [6] the CPR architecture, that includes CTI module, was described, based on the main outcomes of the initial project workshops, whose objective was to understand the circumstances in which the CI operators and authorities found themselves during the COVID- 19 pandemic. In this paper we give a more detailed architecture of CTI module that includes proposed enhancements for “orientation” phase. CTI module is based on open source from MISP Project.

Threat Intelligence Engine (TIE) is the IoC enhancement module that acquires events from the MISP instance and generates a threat score. This score can be divided into two parts: the initial segment is public, as it is founded on open-source data, and it gauges the threat based on diverse metrics like timeliness, trending, and completeness. The second segment, which is private, encompasses these metrics along with the relevance heuristic that factors in CI infrastructure data. To shield against potential information leaks, this segment is encrypted, particularly due to the criticality of its assessment of the infrastructure’s vulnerability against threats. Exposure of this information could have severe consequences, enabling attackers to exploit vulnerabilities

and target the entity. By partitioning the score, we adhere to the principle of sharing data through the public channels (see also Article 29, Cybersecurity information-sharing arrangements and Article 30, Voluntary notification of relevant information, of Network and Information Security – NIS 2 directive [21]), while simultaneously safeguarding the organization's confidential data by encrypting it. Concurrently, we augment the received event's contextual information, procedure referred to as CTI enrichment, described in Figure 2.

The central part of TIE is the Heuristic Engine. It processes API requests containing MISP Events and computes the score by considering information about the critical infrastructure and dynamic data such as events, alerts, vulnerability assessments, temporary conditions etc.

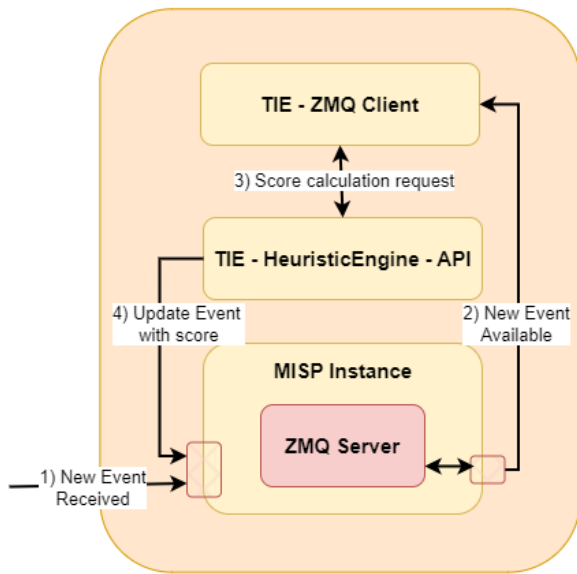


Fig 2. TIE architecture

Subsequently, this score is integrated as an event attribute, updating the MISP Event within the MISP Instance. The second element in TIE's architecture is the ZeroMQ client. A MISP Instance can be configured with a ZeroMQ Server, which has several advantages. The ZMQ Client is established as part of TIE and subscribes to the MISP Instance's ZMQ queue. When a fresh event arrives, the client sends a request to the Heuristic Engine component. This action triggers the execution of heuristic functions that enhances intelligence and ultimately leads to the threat score computation.

TIE currently does not encompass all MISP objects; instead, it concentrates on four specific objects (domain-ip, vulnerability, file, btc-address) considered highly valuable in the realm of threat intelligence and that comprise various attributes, including required and optional ones, which later contribute to the heuristics' calculations and allow the operator to track the main sightings of an attack.

D. CTI enhancements implementation

CTI module, that facilitates secure CTI exchange and enhances cyber-physical risk calculation, has several layers that includes authentication, transport, and privacy/enrichment. The Orchestrator facilitates communication among modules and provides an API for user interaction. The WEB-GUI simplifies tool usage and integrates authentication mechanisms. The new functionality "Source Confidence Calculation" has been implemented and operates in three different modes. Furthermore, admiralty taxonomy is integrated into IoC decay calculation, adjusting IoC significance based on taxonomy classification. Another important feature that has been integrated is the ability to enrich IoC with techniques from the MITRE ATT&CK matrix. Prediction of workforce absenteeism for posterior integration of related intelligence in decisions, such as team assignment, prioritization or calculations of confidence in data sources, is also a new feature from the previous iteration of this component.. It leverages on open-source data, such as anomalies in the number of social network or search engine queries related to health topics, so that CI operators may be able to predict whether their workforce is likely to be unavailable to attend to their workplace due to illness.

E. Assumptions and Constraints

One of the key steps in CPR tool is to identify and map candidate risk mitigations, both to general tactics, techniques and procedures (TTP) from MITRE ATT&CK framework, as well as to security controls specifically applicable to a particular CI. The existing mitigation measure and cyber resiliency controls mappings to ATT&CK techniques are based on engineering analysis rather than on operational experience. In operation settings, risk reduction, and in consequence resilience improvement, will depend on how the controls are specified, how they are implemented, and how the implementation is used. In addition, only the direct effects that a given control are considered. Indirect effects are not considered, while other criteria for the selection of mitigation measures (cost of measure implementation, damage or side-effects) are not considered in this iteration of tool.

F. (Re) Calculation of probabilities, based on CTI

The system was adapted to dynamically update the source's confidence level. The Threat Scoring Model, that contains four heuristics (extensiveness, timeliness, completeness and whitelist overlap) was adjusted, and the output has been used in risk assessment, in order to finetune risk probability. Extensiveness measures how much context an intelligence feed provides to complement additional information. With extensiveness, we assign a higher source confidence to intelligence feeds that give more context per IoC. Timeliness defines how fast an intelligence feed shares its IoCs compared to other feeds. If certain intelligence feed shares the same IoCs later than others, the IoCs could be outdated. Considering the slowness in sharing its IoCs, we will assign less source confidence to it. Completeness defines how much an intelligence feed contributes to the total collection of IoCs. Please note that this score focuses more on the quantity than

on the quality of an intelligence feed, assuming that if it has many IoCs, it indicates the feed is more useful. Whitelist Overlap Score defines how much of the IoCs in an intelligence feed also exist in a trusted whitelist, avoiding the consumption of false positives by the system. These four heuristics are used in a weighted mean to calculate the source confidence of the gathered intelligence feeds. This is the value that will appear in the tag “Source-Score” of the MISP event associated with a feed and is a value exported to risk assessment engine in CPR tool.

An example for risk probability recalculation is threat score parameter “timeliness” that changes perspective of the economic impact. We used historic data about “time to publish” of threat intelligence regarding the new malicious phishing threats with message body that includes COVID-19 information, as well as generic information about average speed of spreading this threat intelligence indicator related to phishing, and a “speed of digestion” in user organization, which is then combined to make calculation about risk probability. In the case that information is shared more than 40 hours before phishing attack on CI entity occurs, which is likely situation, damage reduces significantly, almost to 85 % of what would have been size of damage without the use of threat intelligence.

Finally, we should also mention fine grained policy for CTI sharing. CI organisations are now increasingly aware of the consequences of sharing versus not sharing information, and with new functionalities they can create communities or clusters between entities and employees that work in the same sector or same country, and apply fine grained policy that allows anonymization, encryption, and other operations on threat intelligence data. This is further improving cyber-physical resilience by preparing CI operators against attackers, that use the same TTPs to attack different entities.

IV. VALIDATION

Validation and testing process is scheduled for the summer of 2024 until H2 2025 and will happen in three phases:

- A Proof of Concept will simulate incoming logs from selected applications to identify potential threats to the systems under analysis.
- The tools will be deployed within CI operator to integrate and aggregate logs with the existing SIEM system in the testing environment. Additionally, integrations with the current MISP instance in the user infrastructure may be explored to enhance CTI sharing between departments and organisations.
- The tools will be piloted in the operational environment of CI operator. The tools’ output will be monitored using real data from the applications under analysis, as well as simulated data for vulnerability tests. The CI operator Blue Team might oversee these operations.

In summary, plan is to test and validate this tool on a specific part of its regional critical infrastructure, including VPN Servers and the healthcare application. The testing will progress through phases, beginning with a Proof of Concept,

followed by integration with SIEM and MISP, and concluding with operational environment monitoring.

V. CONCLUSIONS

The COVID-19 pandemic is an example of a temporary situation when CI had to operate with continuously changing conditions. This was reflected in changes of cyber-physical risk assessment, and in consequence, also changes of resilience assessment for CI. Infrastructure parts, employees or assets which were not perceived as critical before the pandemic were revealed to be critical during the pandemic. Due to the temporary conditions, it was not possible to deliver many different supplies for the correct infrastructure operation. The move to remote working highlighted the lack of computers for remote workers, and vulnerabilities with up-to-date configuration or relaxing cybersecurity policies. Most previous approaches describing resilience for critical infrastructure did not mention the essential role of specialists, also termed as “essential workers”. Absenteeism was rarely discussed, as well as the impact on the health or safety of essential workers at their workplace.

Though cyber infrastructures have already been one of the most discussed parts of critical infrastructures prior to the COVID-19 pandemic, its role during pandemics, for example for the remote work or access to critical systems, has also changed. This resulted in re-evaluation of security policies, for example for remote workstation configurations, or higher asset value for some online applications, such as e-health, that must be used by citizens.

Priority or asset value can change and can be very specific for each critical infrastructure. This represents a challenge to keep risk assessment dynamic and in consequence also to increase cyber-physical resilience.

We started from investigation of these challenges with mappings of different static and dynamic inputs from internal “observation” sources that are relevant for cyber-physical risk models, usually a combination of events that can be collected through existing cyber tools, and information that can be statically provided by the CI operator. In addition, external sources of information, such as threat intelligence, were considered and modelled to make an impact on risk assessment. These external sources and events were also enhanced with threat score, as well as enriched with strategic information about adversary TTPs through “orientation” phase, and different heuristics of calculated score were considered for the re-calculation and adjustment of risk assessment.

Our approach shows that some temporary changes need to be modelled in advance and that collaboration through tools such as threat intelligence sharing, as well as dynamic enhancements and adaptations to temporary conditions, is essential in adjustment to temporary conditions.

In summary, we show how sharing and enhancing threat intelligence can have several impacts on CI resilience, including:

- Increased awareness about dynamic changes and temporary conditions,

- Improved detection capabilities (e.g. decreased time to develop new rules for intrusion detection systems or multi-factor authentication),
- Enhanced risk assessment by adjusting the potential impact of attacks, as well as the likelihood of occurrence.
- Informed decision-making to allocate scarce resources more effectively, prioritize security controls and countermeasures, and invest in relevant security technologies and training programs.

There is still work to do in the SUNRISE project, including validation by the end users and increasing the complexity of system under observation, according to the definition of complex system [22]. Balance might need to be found between the system dynamics, organisational priorities, and behavioural aspects. We need re-assessment every time there is a change of risk inputs or indicators, change of dependencies between them, or if a gap between strategic/tactical level decisions and the operational cybersecurity management (monitoring, detection, and operational response) has been detected. This might lead to “assessment fatigue” or even congestion in decision making. Finetuning might be needed to reach these balances and work on trade-offs between dealing with dynamicity and tool practicality.

By combining sensing (“observation”) of external environment (incidents from SIEM or IDS tools, threats from CTI platforms, vulnerability from scanners, abnormal behaviour events, etc.), with a cognitive process of “orientation” (including threat score calculation or mapping to TTPs), we move towards cyber-physical resilience which is having ability to anticipate unknown faults or incidents. In this direction, we think that new research is needed in the areas related to dynamic and adaptable threat intelligence sharing, probably as a part of a different research project.

ACKNOWLEDGMENT

This work has been supported by the European Commission through SUNRISE project (grant no. 101073821) under the Horizon Europe research programme.

REFERENCES

- [1] Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, NIST Special Publication 800-160, Volume 2, Revision 1, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v2r1.pdf>
- [2] National definitions of Cyber resilience, CIPedia: https://websites.fraunhofer.de/CIPedia/index.php/Cyber_Resilience
- [3] Birkie, Seyoum Eshetu & Trucco, Paolo & Kaulio, Matti. (2014). Disentangling core functions of operational resilience: a critical review of extant literature. *Int. J. of Supply Chain and Operations Resilience*. 1. 76-103. 10.1504/IJSCOR.2014.065461.
- [4] Council directive 2008/114/EC of 8 December 2008 on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection, <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ.L:2008:345:0075:0082:EN:PDF>
- [5] Directive (EU) 2022/2557 of the European Parliament and of the Council of 14 December 2022 on the resilience of critical entities and repealing Council Directive 2008/114/EC, <https://eur-lex.europa.eu/eli/dir/2022/2557/oj>
- [6] P. de Juan Fidalgo, A. Pasic, J. M. Del Álamo, R. Touris and A. Álvarez, "TERME: a cyber-physical resilience toolset for risk assessment," 2023 JNIC Cybersecurity Conference (JNIC), Vigo, Spain, 2023, pp. 1-6, doi: 10.23919/JNIC58574.2023.10205687.
- [7] ENISA technical report, Measurement Frameworks and Metrics for Resilient Networks and Services: Technical report, February 2011
- [8] ENISA report Ontologies and Taxonomies for Resilience , 2011 https://www.enisa.europa.eu/publications/ontology_taxonomies
- [9] Eisenberg, Daniel, Plourde, Kenton, Seager, Thomas, Allen, Julia & Kott, Alexander. (2013). Resilience metrics for cyber systems. *Environment Systems and Decisions*. 33. 10.1007/s10669-013-9485-y.
- [10] Gustavo González-Granadillo, Mario Faiella, Ibéria Medeiros, Rui Azevedo, Susana González-Zarzosa, ETIP: An Enriched Threat Intelligence Platform for improving OSINT correlation, analysis, visualization and sharing capabilities, *Journal of Information Security and Applications*, Volume 58,2021, 102715, ISSN 2214-2126, <https://doi.org/10.1016/j.jisa.2020.102715>.
- [11] Faiella, M.; Gonzalez-Granadillo, G.; Medeiros, I.; Azevedo, R. and Gonzalez-Zarzosa, S. (2019). Enriching Threat Intelligence Platforms Capabilities. In *Proceedings of the 16th International Joint Conference on e-Business and Telecommunications - SECURITY*; ISBN 978-989-758-378-0; ISSN 2184-3236, SciTePress, pages 37-48. DOI: 10.5220/0007830400370048
- [12] G. Gonzalez-Granadillo, M. Faiella, I. Medeiros, R. Azevedo and S. Gonzalez-Zarzosa, "Enhancing Information Sharing and Visualization Capabilities in Security Data Analytic Platforms," 2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W), Portland, OR, USA, 2019, pp. 1-8, doi: 10.1109/DSN-W.2019.00009.
- [13] World Health Organisation initiative: <https://www.who.int/initiatives/eios>
- [14] Abbas H, Tahoun MM, Aboushady AT, Khalifa A, Corpuz A, Nabeth P. Usage of social media in epidemic intelligence activities in the WHO, Regional Office for the Eastern Mediterranean. *BMJ Glob Health*. 2022 Jun;7(Suppl 4):e008759. doi: 10.1136/bmjgh-2022-008759. PMID: 35764352; PMCID: PMC9240825.
- [15] European Centre for Disease Prevention and Control web page: <https://www.ecdc.europa.eu/en/information-social-media-monitoring-epidemic-intelligence-purposes>
- [16] A. K. Daou, F. Li and S. Shiaeles, "A Cost-Efficient Threat Intelligence Platform Powered by Crowdsourced OSINT," 2023 IEEE International Conference on Cyber Security and Resilience (CSR), Venice, Italy, 2023, pp. 48-53, doi: 10.1109/CSR57506.2023.10225008.
- [17] Anirban Mukhopadhyay, Sukrit Venkatagiri, and Kurt Luther. 2024. OSINT Research Studios: A Flexible Crowdsourcing Framework to Scale Up Open Source Intelligence Investigations. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 105 (April 2024), 38 pages. <https://doi.org/10.1145/3637382>
- [18] MISP platform to gain situational awareness in regards to the COVID-19 situation: <https://www.misp-project.org/covid-19-misp/>
- [19] MISP taxonomy for COVID-19: <https://github.com/MISP/misp-taxonomies/blob/main/pandemic/machinetag.json>
- [20] Toolkit for Mapping of the MISP for SRH and its Adaptation for Preparedness and Response to COVID-19 and Other Pandemics and Major Outbreaks, <https://iawg.net/resources/toolkit-for-mapping-of-the-misp-for-srh-and-its-adaptation-for-preparedness-and-response-to-covid-19-and-other-pandemics-and-major-outbreaks>
- [21] Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union, also known as NIS2 Directive, <https://eur-lex.europa.eu/eli/dir/2022/2555>
- [22] Jan Žižka, Bruno Rossi, Tomáš Pitner, Towards a Definition of Complex Software System, Position Papers of the 18th Conference on Computer Science and Intelligence Systems, M. Ganzha, L. Maciaszek, M. Paprzycki, D. Ślęzak (eds). ACSIS, Vol. 36, pages 119–126 (2023), DOI: <http://dx.doi.org/10.15439/2023F2898>

Hospital Patient Distribution After Earthquake

Stefka Fidanova

Institute of Information and
Communication Technology
Bulgarian Academy of Sciences
Sofia, Bulgaria
E-mail: stefka.fidanova@iict.bas.bg

Veselin Ivanov

Faculty of Medicine,
Trakia University,
Stara Zagora, Bulgaria
E-mail: veskoasenov@abv.bg

Leoneed Kirilov

Institute of Information and
Communication Technology
Bulgarian Academy of Sciences
Sofia, Bulgaria
E-mail: l_kirilov_8@abv.bg

Maria Ganzha

System Research Institute
Polish Academy of Sciences
Warsaw, Poland
E-mail: maria.ganzha@ibspan.waw.pl

Abstract—The correct organization of medical assistance after the occurrence of a major disaster is very important for saving the lives of the victims. Earthquakes are natural phenomena/disasters in which there are many victims. The timely provision of medical assistance to the injured is an important element of their service. It is good to divide them into types of injuries and severity of injuries. Thus, the medical teams will be prepared for how many people need outpatient treatment and how many need hospital treatment. Rapid distribution of victims to hospitals according to their injuries can reduce the number of deaths and people with serious consequences. In this article, we present a breakdown of the injured by hospitals and medical facilities near the earthquake site. The type of injuries and the capacity and equipment of hospital facilities are taken into account.

Index Terms—Patient flow optimization, Earthquake, Distribution to hospitals

I. INTRODUCTION

WHEN a disaster situation occurs, such as a strong earthquake, the organization of medical assistance is of great importance for the rescue of the victims. The availability of information about the hospitals in the area, their equipment and capacity can be used for distribution of the injured. Their ability to provide adequate assistance to the wounded should be known in advance [1]. Earthquake zone maps exist for some areas. They are based on statistical and geological data and give the probability of an earthquake with a large magnitude [2]. The occurrence of a disaster situation is something of an emergency for hospitals, and for this there must be advance arrangements for the reception of victims of a certain type. The preliminary organization is to predict, calculate what number of victims there would be and of what type, knowing the type of construction, the occupants of each building and the type of heating they use. Two types of assessments can be made, one for the heating season and one outside the heating season. During the heating season, the likelihood of fires and people with burns is much greater than outside

the heating season [3]. In the event of a disaster, hospitals must very quickly reorganize their organization for receiving and handling incoming patients, because they will receive a much larger number of patients than usually a certain type of patient. This is called the hospital's ability to respond and is determined by the hospital's normal capacity and the ability to expand [4].

$$RespondCapability = Planning \times Capacity$$

where the *Plannin* is how much the capacity can be expended. Planning shows what the health facility's potential is if it uses all its available resources, such as staff, equipment, supplies.

The need for disaster medical preparedness planning and timely response was clearly seen in the corona virus pandemic [5]. This includes various components such as space, infrastructure, staff (medical and non-medical), medical and non-medical equipment. Although essential, planning and preparing for disaster response is often neglected. The reason may be the heavy workload with the daily sick [6].

The World Health Organization gives prescriptions for action in various disaster and pandemic situations, but they are too general and rather advisory, which makes it difficult to apply them directly [7]. The United States and the European Union also have developed prescriptions and disaster response plans. In some European countries, emergency departments are required by law to implement these prescriptions and plans, but they are also too general [8]. These are rather action plans for a particularly large flow of patients and are not entirely related to the disaster situation.

Giving first aid to the injured at the scene of the disaster, classifying them according to types of injuries and their severity, sending the seriously injured to hospital facilities is of great importance. This requires prior organization and

material preparation, which is complex and should not be underestimated.

Therefore, mathematical models are made to be able to assess the disaster situation, the number and types of victims and what kind of help they need. A model is needed to be able to assess the specific situation and propose a solution for the disaster area. Thus, in different regions, the situation and the decisions to be taken will be different, tailored to the specific situation. This development aims to optimize the distribution of earthquake victims by hospital. Good organization helps to save lives, reduce and even prevent panic among the population.

The rest of the paper is organized as follows. Section 2 is devoted to literature review. Section 3 provides an assessment of the victims by type and number. In Section 4, an algorithm for the distribution of the injured by hospital facilities is presented. Section 5 provides a conclusion and direction for future research.

II. LITERATURE REVIEW

Advance planning, procurement and overall organization of rescue activities and medical care after an earthquake is extremely important to save people and prevent further consequences. This can be achieved by creating models showing the types of damage that would occur and the potential casualties. The types of construction in the considered region, the occupants of the individual buildings, the health facilities near the area of destruction and their capabilities, the expected magnitude of the possible earthquake can serve as input data. As a result, such models can show the type and number of victims, the medical assistance they need, the distribution of victims by health facilities, the necessary minimum of medicines and sanitary material that hospitals in the area must maintain.

Most of the existing models are aimed at rapid reconstruction of a hospital facility in case of a strong increase in the flow of patients, many times greater than normal [9]. A disadvantage of this approach is that it requires too much time and often lacks realism [10]. There are smaller patterns of patient flow that apply to specific departments in the hospital [11]. Other models are more comprehensive and consider the hospital and the flow of all patients as a whole [12]. These models are too general and do not reflect the reason for the increased admission of patients and their specificity.

There are models that are tailored to the disaster event, but again they are geared towards handling the increased number of patients. A model is proposed in [13] relating to increased patient flow due to an influenza epidemic. In [14], the model refers to patient care after a major fire, but again it is a single hospital. In [15], the model is for patient flow after an explosion.

Our proposed model targets patients at the scene of the disaster. First, an assessment is made of the possible victims and how many of them need hospital treatment. Hospital capacity information is used, starting with the nearest, until

all those in need are served. The other models are primarily aimed at serving the growing number of patients.

III. ASSESSMENT OF PATIENTS

There are a variety of models for assessing earthquake victims. The devastating earthquake consequences for the environment are a major cause for the occurrence of numerous medical emergencies. There are usually large numbers of permanent (deaths) and temporary (injured) medical losses. The types of injuries are diverse in type and structure. The direct impact trauma factor is the most common in cases of earthquake, as the direct impact is usually caused by airborne debris and falling objects from demolishing buildings, thus inflicting mostly traumatic injuries - see Clark (2018) [16]. In cases of secondary disaster area occurrences, it is possible to have impact factors like radiation; toxic substance release; thermal impact, biological impact, etc., which can cause radiation damage; acute poisonings; thermal injuries; outbreaks of infectious diseases; cases of drowning. The negative effects of the psychological stress factor are inevitable and they cause acute neuropsychological conditions in vast numbers of people - see Farooqui et al. (2017) [17], Todorova et al. (2020) [18], Etova (2021) [19].

In Tirkolae et al. (2020) [20] a robust bi-objective mixed-integer linear programming model to allocate disaster rescue units is proposed. The authors apply the model to a real case study for Mazandaran province in Iran. The same problem for resource allocation for emergency response is studied in Fiedrich et al. (2000) [21]. They develop dynamic optimization model and a method for solving it. The goal is to present a schedule for optimizing of the available technical resources. Robust Model for Logistics Management (RMLM) is proposed in Najafi et al. (2012) [22]. The model is multi-objective, multi-mode, multi-commodity, and multi-period. It is proposed to manage the logistics of commodities and injured people in the case of earthquake. A three objective mixed integer stochastic model for locations of storage areas for shelters pre-earthquake and distribution of shelters is proposed in Yenice and Samanlioglu (2020) [23]. The authors consider four event scenarios according to two different earthquake scenario likelihoods. The model was applied to Kadikoy municipality of Istanbul, Turkey. Dawei et al. (2015) [24] the problem about vehicle scheduling in the medicine dispatching process is studied in order to minimize the total transportation time under several types of vehicles. The problem is solved by genetic algorithm.

Transient modelling with quadratic regression analysis for simulation of hospital operations in emergency situations is studied in Paul et al. (2006) [25]. A double exponential function is used to model the transient waiting time. De Boer and Debacker (2006) [26] determine the medical resources for disasters in the Netherlands. They study the medical rescue capacity, the medical transport capacity and the hospital treatment capacity using medical severity index model. Several cases are considered under different assumptions of severity. A stochastic Petri net is used for modelling and optimizing

the emergency medical rescue (EMR) process in Sun et al. (2021) [27]. The approach is tested with the data of the 2008 Wenchuan earthquake. Ghasemi et al. (2020) [28] propose a stochastic multi-objective mixed-integer model for logistic distribution and evacuation planning during an earthquake. The model is converted into deterministic one and then solved by means of genetic algorithm NSGA-II. The model is tested for a probable earthquake in Tehran.

In [3] is proposed a model, based on Generalized Nets, for processes occurring during an earthquake. With the help of this model, damage and casualties can be estimated when we know the types of construction, its earthquake resistance, the number of occupants of each building and the type of heating.

The injured will be divided into lightly injured and severely injured. Minor injuries can receive medical attention on the spot and do not require hospitalization. The seriously injured are people who will receive first aid on the spot and then be directed to hospital treatment. We divide the severely injured into two main groups. Victims with fractures and victims with burns.

Let the number of residents of a given building be N . If the earthquake resistance of the building is less than the strength of the earthquake, then there will be major damage to the building and we assume that the number of occupants with fractures will be F , where:

$$F = 0.25 * N$$

If solid fuel heating is used in this building, local fires are likely to occur and the number of people with burns will be B , where:

$$B = 0.1 * N$$

In case of strong earthquakes with a magnitude above 7, if there is a gas installation in the building, there is a high probability of damage to the installation and gas leakage. This could result in a serious fire and people with severe burns. In this case, we assume that people with severe burns will be B , where:

$$B = 0.5 * N$$

All these victims need hospital treatment. Of course, in the preparation of the emergency teams, people with minor injuries and minor burns, who will not need hospital treatment, but sanitary materials will be used for them, should also be taken into account. For people with serious injuries, sanitary materials will also be used in providing first aid on the spot before being taken to a hospital facility. The percentage of casualties out of the total number of occupants was determined in consultation with disaster medicine specialists. These percentages are input parameters to the model and can be changed and refined.

IV. PATIENT DISTRIBUTION TO HOSPITALS

The model of the consequences of an earthquake, which we have done in our previous development, estimates the types of damage from a given area [3]. They are grouped into two types of injuries, fractures and burns. Each type has two main subtypes, heavy and light. Severe ones require hospital treatment, while mild ones can be treated on site and patients referred to accommodation centers. Of interest to this article are the victims who need hospital treatment and their distribution by hospital.

Let there be n number of hospitals within a radius of 50 km from the epicenter of the earthquake $\{H_1, H_2, H_3, \dots, H_n\}$, where H_1 is the closest hospital to the epicenter and H_n is the most distant hospital to the epicenter. For each hospital, there is a vector with information about it.

$$H_i = (d_i, cf_i, cf1_i, cb_i, cb1_i)$$

where d_i is a distance from the earthquake epicenter to the hospital, cf_i is the capacity of the trauma department and $cf1_i$ is the capacity of possible extension of trauma department; cb_i is the capacity of the hospital's burn department and $cb1_i$ is the possible extension of the burn department.

If the hospitals within a radius of 50 km do not have sufficient capacity (including extended capacity) to accommodate the expected number of victims, new hospitals are added, increasing the radius by 10 km. We increase the radius until enough hospitals are involved to accommodate all the seriously injured.

Some of the smaller hospitals only have a trauma unit and no burn unit. Then the value for the capacity of the burn unit is $cb_i = 0$ and for extension of the burn unit is $cb1_i = 0$.

When drawing up a plan for the distribution of injured people, it is first assessed which of the available hospitals have an earthquake resistance equal to or greater than the magnitude of the expected earthquake. This means that the hospital will not have significant damage and will be able to function. If the earthquake resistance of the hospital building is less, then we set its capacity to be $cf_i = 0$, $cf1_i = 0$, $cb_i = 0$ and $cb1_i = 0$. Thus, this hospital will not enter the casualty distribution plans. The goal is that the total time for accommodation of the victims is minimal.

We offer the following algorithm for the distribution of the injured by hospitals:

- 1 The victims are distributed according to the severity of the injury;
- 2 $i = 1$;
- 3 If $cf_i > 0$ $vf_i = \min\{cf_i, F\}$; The distribution of fracture victims starts from the closest hospital to the epicenter with capacity $cf_i > 0$, vf_i are fracture victims distributed to H_i , F are not distributed fracture victims;
- 4 After the places are filled, this hospital is removed from the list by placing $cf_i = 0$;
- 5 $F = F - vf_i$;
- 6 $i = i + 1$;

- 7 If $i \leq n$ and there are unallocated casualties ($F > 0$), return to position 3;
- 8 If $i > n$ and $F > 0$, return to position 2 using $cf1_i$ instead of cf_i , otherwise we move to position 9;
- 9 All the injured have been transferred to appropriate hospitals.

A similar scheme is also used for the distribution of victims with severe burns.

The radius of the hospitals in which victims with fractures are distributed may be different from the radius of hospitals in which the victims with severe burns are distributed. This is because, on the one hand, the victims of earthquake with fractures, especially outside the heating season, are much more than the victims with burns. On the other hand, most hospitals for general treatment, including small ones, have trauma departments. This is not the case with burn departments.

When determining the capacity of a department, the possibility of placing additional beds and serving additional sick/injured patients is also taken into account. The additional beds/casualties represent additional overburdening of the hospital and, from there, a possible decrease in the quality of the service. Additional beds are filled when all hospitals in the area are at capacity and there are unserved casualties. The same algorithm is applied for extensions of the hospitals. In the proposed algorithm, a balance is sought between the time to take the patients to a hospital and the quality of the health service offered.

V. COMPUTATIONAL EXAMPLES

We have prepared the following example to show how our proposed algorithm works.

The example consists of residential region with: 8 high apartment buildings with 210 inhabitants each; 42 apartment buildings up to 3 floors with 36 inhabitants each and 500 family houses with 4 inhabitants each. The number of all inhabitants of the region is 5192. There are various types of heating.

Let there be one hospital in the settlement. This hospital has a trauma department with a capacity of 50 beds and can be expanded to 150 beds. The hospital does not have a burn unit. At a distance of 10 km there is a hospital with a trauma department with a capacity of 100 beds and the possibility of expansion to 300, as well as a burn unit with a capacity of 10 beds and the possibility of expansion to 20. At a distance of 20 km. there is a hospital with a trauma department with a capacity of 50 beds, expandable to 200, without a burn department. At a distance of 30 km. there is a hospital with a trauma department with a capacity of 300 beds, expandable to 800, and a burn unit with a capacity of 20 beds, expandable to 80.

When the magnitude of the earthquake is up to 7 there will be no died people. The people with minor injures will be 463, with fractures they will be 200 and 40 with minor burns, because some of the houses use solid fuel heating and it can cause local fires.

In this case, there are 200 people with fractures, 50 of them will be admitted to the first hospital, 100 people will be admitted to the hospital 10 km away and the remaining 50 will be admitted to the extra beds in the first hospital. No one suffered serious burns, only minor burns were reported and they could be treated at the scene by paramedics or emergency responders.

When the magnitude of the earthquake is 8 and more there will be 689 died people. The people with minor injures will be 1919, with fractures they will be 1246. There is a possibility for explosion of gas installation, therefore there will be 15 peoples with hard burns and 206 with minor burns.

In this case, there are 1,246 people with fractures. 50 of them will be accommodated in the first hospital. In the hospital 10 km away, 100 people will be accommodated. In the hospital 20 km away, 50 people will be accommodated and 300 people will be accommodated in the hospital 30 km away.

The remaining 646 people will be accommodated on the extra beds as follows. 150 people will be accommodated in the first hospital. In the hospital 10 km away, 300 people will be accommodated. The remaining 196 people will be accommodated in the hospital 20 km away. In the hospital 30 km away, no additional beds will be accommodated.

We have 15 victims with severe burns who need hospital treatment. There are burn departments in the hospital 10 km away and in the hospital 30 km away. The hospital 10 km away has a capacity of 10 beds with the possibility of expansion by another 20. All burn victims will be taken to this hospital because the number of victims exceeds the normal capacity of the hospital by only 5 people and the expanded capacity is not exceeded.

VI. CONCLUSION

In this article, we have proposed a model for the distribution of earthquake victims in hospital facilities. The type and number of victims, the equipment of the hospitals and their capacity, their ability to reorganize their activities were taken into account. The distance of the hospital to the scene of the accident was also taken into account. Patients are sent to the nearest possible hospital that can provide them with the necessary care until its capacity is filled. The goal is to optimize medical care.

The purpose of the model is a preliminary assessment of the injured and their need for medical care. On the basis of this assessment, a preliminary plan can be drawn up for the distribution of the injured by hospitals and medical centers. The model can serve for the training and preparation of medical teams, as well as for the material security of hospitals.

As future work, the model will be further developed and refined. On its basis, models could be created for medical assistance in other types of disaster situations.

ACKNOWLEDGMENT

The work is supported by National Scientific Fund of Bulgaria under the grant DFNI KP-06-N52/5 and by the Polish-Bulgarian collaborative grant "Practical aspects for scientific computing".

REFERENCES

- [1] Liguori N., Tarque N., Bambaren C., Spacone E., Viveen W., de Filippo G., *Hospital treatment capacity in case of seismic scenario in the Lima Metropolitan area, Peru*, Int. J. Disaster Risk Reduc. Vol. 38 Elsevier, 2019, 101196.
- [2] Toner E., Schoch-Spana M., Waldhorn R., Shearer M., Inglesby T., *A Framework for Healthcare Disaster Resilience: A View to the Future*, 2018, <http://www.centerforhealthsecurity.org/our-work/publications/a-framework-for-healthcare-disaster-resilience-a-view-to-the-future>.
- [3] Fidanova, S., Atanassov, K., Kirilov, L., Slavova V., Ivanov, V.: Generalized Net Model for the Consequences of Earthquake. Lecture Notes in Networks and Systems, 658, 281-292 ISBN:978-3-031-31068-3, ISSN:2367-3370, 2023
- [4] Hick J.L., Barbera J.A., Kelen G.D., *Refining surge capacity: conventional, contingency, and crisis capacity*, Disaster Med. Public Health Prep. Vol. 3, 2009, S59-S67.
- [5] Peleg K., Bodas M., Hertelendy A.J., Kirsch T.D., *The COVID-19 pandemic challenge to the All-Hazards Approach for disaster planning*, Int. J. Disaster Risk Reduc. Vol. 55, 2021, 102103.
- [6] Lim H.W., Li Z., Fang D., *Impact of management, leadership, and group integration on the hospital response readiness for earthquakes*, Int. J. Disaster Risk Reduc. Vol. 48, 2020, 101586.
- [7] World Health Organization, *Surge Planning Tools* (n.d.), <https://www.euro.who.int/en/health-topics/Health-systems/pages/strengthening-the-health-system-response-to-covid-19/surge-planning-tools>.
- [8] Gazzetta Ufficiale della Repubblica Italiana, D.lg.s 9 aprile, Testo Unico Sulla Salute e Sulla Sicurezza sul Lavoro, 2008. Italy, 2008, [https://www.lavoro.gov.it/documenti-e-norme/studie-statistiche/Documents/Testo Unico sulla Salute e Sicurezza sul Lavoro/Testo-Unico-81-08-Edizione-Giugno 2016.pdf](https://www.lavoro.gov.it/documenti-e-norme/studie-statistiche/Documents/Testo%20Unico%20sulla%20Salute%20e%20Sicurezza%20sul%20Lavoro/Testo-Unico-81-08-Edizione-Giugno%202016.pdf).
- [9] Farra S.L., Gneuh M., Hodgson E., Kawosa B., Miller E.T., Simon A., Timm N., Hausfeld J., *Comparative cost of virtual reality training and live exercises for training hospital workers for evacuation*, Comput. Inf. Nurs. 37, 2019, 446-454.
- [10] Verheul M.L., Duckers M.L.A., Visser B.B., Beerens R.J., Bierens J.J., *Disaster exercises to prepare hospitals for mass-casualty incidents: does it contribute to preparedness or is it ritualism?*, Disaster Med 33, 2018 387-393.
- [11] Rohleder T.R., P. Lewkonja P., D.P. Bischak, P. Duffy, R. Hendijani, *Using simulation modeling to improve patient flow at an outpatient orthopedic clinic*, Health Care Manag. Sci. 14 (2011) 135-145, <https://doi.org/10.1007/s10729-010-9145-4>.
- [12] TariVerdi M., Miller-Hooks E., Kirsch T., *Strategies for improved hospital response to mass casualty incidents*, Disaster Med. Public Health Prep. 12, 2018, 778-790.
- [13] Laskowski M., Demianyk B.C.P., Witt J., Mukhi S.N., Friesen M.R., McLeod R.D., *Agent-based modeling of the spread of influenza-like illness in an emergency department: a simulation study*, IEEE Trans. Inf. Technol. Biomed. a Publ. IEEE Eng. Med. Biol. Soc. 15, 2011 877-889.
- [14] Abir M., Davis M.M., Sankar P., Wong A.C., Wang S.C., *Design of a model to predict surge capacity bottlenecks for burn mass casualties at a large academic medical center*, Prehospital Disaster Med. 28, 2013, 23-32.
- [15] Hirshberg A., Scott B.G., Granchi T., Wall M.J.J., Mattox K.L., Stein M., *How does casualty load affect trauma care in urban bombing incidents? A quantitative analysis*, J. Trauma Acute Care Surg. 58 2005.
- [16] Clark, K. R.: *Imaging Earthquake-related Injuries*. Radiol. Technol. 2018 Mar. 89(4), 351-367. (2018) PMID: 29691346.
- [17] Farooqui, M., Quadri, S.A., Suriya, S., Khan, M. A., Ovais, M., Sohail, Z., Shoaib, S., Tohid, H., Hassan, M.: *Posttraumatic stress disorder: a serious post-earthquake complication*. Trends Psychiatry Psychother. 2017 Apr-Jun. 39(2), 135-143 (2017) doi: 10.1590/2237-6089-2016-0029. PMID: 28700042.
- [18] Todorova, D., Mihaylova, Tsv., Etova, R.: *Social psychological support for disasters*. Health Policy and Management. vol. 20, Extraordinary issue. 169-172. ISSN 1313-4981 (2020)
- [19] Etova, R.: *Disasters and stress*. Scientific works of the Union of Scientists in Bulgaria Plovdiv. series G. Medicine, Pharmacy and Dental medicine. vol. 26, 190-193. ISSN 1311- 9427 (Print), ISSN 2534-9392 (On-line) (2021)
- [20] Tirkolaei, E. B., Aydın, N. S., Ranjbar-Bourani, M., Weber, G.-W.: *A robust bi-objective mathematical model for disaster rescue units allocation and scheduling with learning effect*. Computers & Industrial Engineering. 2020. <https://doi.org/10.1016/j.cie.2020.106790>
- [21] Fiedrich, F., Gehbauer, F., Rickers, U.: *Optimized resource allocation for emergency response after earthquake disasters*. Safety Science 35 (2000) 41-57.
- [22] Najafi, M., Eshghi, K., Dullaert, W.: *A multi-objective robust optimization model for logistics planning in the earthquake response phase*. Transportation Research Part E 49 (2013) 217-249. <http://dx.doi.org/10.1016/j.tre.2012.09.001>
- [23] Yenice, Z. D., Samanlıoğlu, F.: *A Multi-Objective Stochastic Model for an Earthquake Relief Network*. Journal of Advanced Transportation. vol. 2020, Article ID 1910632 (2020). <https://doi.org/10.1155/2020/1910632>
- [24] Dawei, L., Yingying, X., Li, W.: *Vehicle Scheduling of the Emergency Medicines in the Early Period After Earthquake Disaster*. The Open Cybernetics & Systemics Journal. 9, 1329-1333 (2015).
- [25] Paul, J.A., George, S.K., Yi, P., Lin, L.: *Transient modeling in simulation of hospital operations for emergency response*. Prehospital and Disaster Medicine 2006, 21(4), 223-236 (2006).
- [26] de Boer, J., Debacker, M.: (2006) *Quantifying medical disaster management*. International Journal of Disaster Medicine. 000: 1-5. (2006). DOI: 10.1080/15031430600975569
- [27] Sun, H., Liu, J., Han, Z., Jiang, J.: *Stochastic Petri Net Based Modeling of Emergency Medical Rescue Processes During Earthquakes*. J. Syst. Sci. Complex. 34, 1063-1086 (2021). DOI: 10.1007/s11424-020-9139-3
- [28] Ghasemi, P., Khalili-Damghani, K., Hafezalkotob, A., Raissi, S.: *Stochastic optimization model for distribution and evacuation planning (A case study of Tehran earthquake)*. Socio-Economic Planning Sciences. 71 (2020) 100745. <https://doi.org/10.1016/j.seps.2019.100745>

Impact of Spelling and Editing Correctness on Detection of LLM-Generated Emails

Paweł Gryka, Kacper Gradoń, Marek Kozłowski, Miłosz Kutyla, Artur Janicki

0009-0002-8505-2098

0000-0003-0750-8678

0000-0002-6313-8387

0009-0002-0947-8986

0000-0002-9937-4402

Warsaw University of Technology

ul. Nowowiejska 15/19, 00-665 Warsaw, Poland

Email: {Pawel.Gryka.stud, Kacper.Gradon, Marek.Kozlowski, Milosz.Kutyla.stud, Artur.Janicki}@pw.edu.pl

Abstract—In this paper, we investigated the impact of spelling and editing correctness on the accuracy of detection if an email was written by a human or if it was generated by a language model. As a dataset, we used a combination of publicly available email datasets with our in-house data, with over 10k emails in total. Then, we generated their “copies” using large language models (LLMs) with specific prompts. As a classifier, we used random forest, which yielded the best results in previous experiments. For English emails, we found a slight decrease in evaluation metrics if error-related features were excluded. However, for the Polish emails, the differences were more significant, indicating a decline in prediction quality by around 2% relative. The results suggest that the proposed detection method can be equally effective for English even if spelling- and grammar-checking tools are used. As for Polish, to compensate for error-related features, additional measures have to be undertaken.

I. INTRODUCTION

ONE of the most serious problems associated with the developments in Information Technologies and their public availability today is the detection of content generated by Artificial Intelligence (AI). Apart from the substantial benefits introduced by public AI applications (especially in such fields as medical imaging diagnostics, data analysis, or automated translation), there is also a set of undeniable challenges caused by the constantly increasing difficulty in distinguishing between human and machine-generated text, images, video, and audio.

These problems are highlighted by transnational law-enforcement institutions, such as Europol [1], the intelligence and national security community [2] or public health policymakers and researchers [3], who emphasize the threats related to the application of the Generative Artificial Intelligence (GAI) for the creation of disinformation, sophisticated scams, social engineering and political manipulation. The GAI-related challenges do not have to be linked to high-profile security domains only. An important and worrying abuse of technology can also be seen in the academic world, where GAI brings

Research was funded by the Warsaw University of Technology within the Excellence Initiative: Research University (IDUB) programme.

plagiarism and academic dishonesty to an entirely new level. Early detection and flagging of high-quality, language-agnostic content produced by AI tools require urgent research and development efforts and the creation of high-quality detection tools.

This paper is a continuation of our efforts to recognize LLM-generated texts, initially focused on email messages. The preliminary results of our experiments have been described in our previous work [4]. The detection results yielded F1-scores of almost 0.98 for English and over 0.92 for Polish. It turned out that the detection algorithm strongly relied on sentence statistics, such as the average word and sentence length, as well as on typographical and orthographic (spelling) imperfections. However, those experiments did not consider that the analyzed text might have undergone spelling and grammar checks. In this study, we would like to check what the actual impact of those errors on detection accuracy is.

Our paper is structured as follows: first, in Section II, we summarize related work in this field. In Section III, we describe our approach. Section IV presents the methodology of our experiments. Results are shown and discussed in Section V, followed by conclusions in Section VI.

II. RELATED WORK

Human communication is increasingly flooded by AI-generated texts. LLMs suggest words and paragraphs or produce entire essays across chat, email, and social media. Therefore, there is a huge need for an effective method of detecting LLM-generated texts (LLMGT).

Several approaches for LLMGT detection have been suggested and explored. Some researchers proposed watermarking or registering AI-generated content. The main idea of these approaches is that any organization developing a foundation model intended for public use must demonstrate a reliable detection mechanism for the content it generates as a condition of its public release. Knott et al. [5] proposed using watermarking as a solution for detecting LLMGT. The authors claim that searching for watermarks can be very effective. Another approach [6] relies on retrieving semantically-

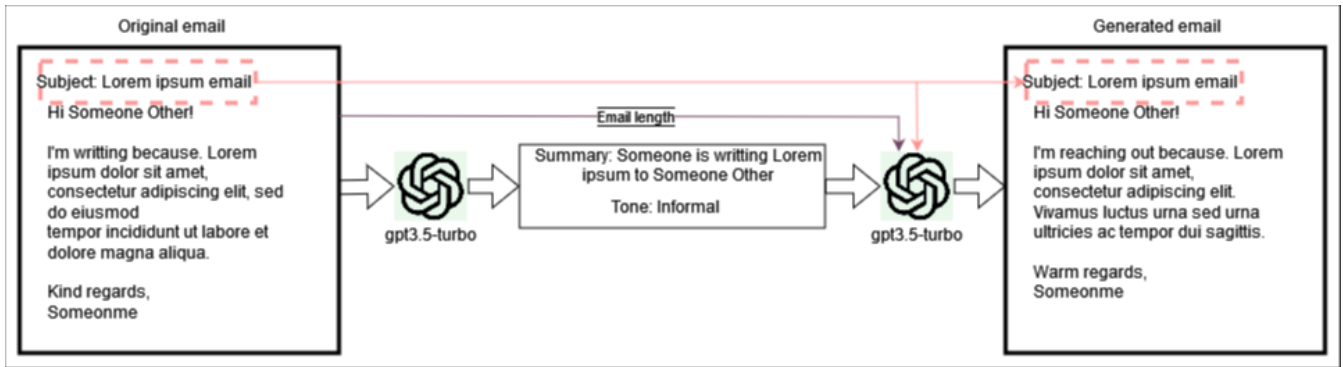


Fig. 1. Process of creating LLM-generated version of email.

similar generations from a huge database with language model historical outputs.

Many researchers use feature-based detection, which seems to be the most straightforward. It consists of extracting various text characteristics in hopes of finding differences between human-written and generated texts. Then, those characteristics are used to train a machine learning (ML) model that will be able to analyze them and produce classification results. Many possible characteristics can be measured; those include stylometry features like word frequency and sentence structure, LLM-specific features like perplexity [7] (measure representing how “surprised” the specific LLM model is when seeing the given text; human-written texts are usually characterized with higher perplexity scores than LLM generated ones) and burstiness [8] (metric based on words’ distribution and the variance of sentence length in the text; humans tend to write with a less consistent style than an LLM might). Their effectiveness was demonstrated by Cingillioglu [9], who used linguistic, semantic, and stylistic features to train a support vector machine (SVM) classifier and got over 92% accuracy when detecting generated essays. This approach was also used by Fröhling and Zubiaga [10], who observed that AI-generated texts exhibit 1) lack of syntactic and lexical diversity, 2) repetitiveness (tendency to overuse frequent words), 3) lack of coherence, and 4) lack of clear purpose or focus.

Another group of methods uses LLMs themselves to detect LLM-generated texts. The first set of methods applies LLMs as they are, i.e., without any training steps. Shi et al. [11] introduced Proxy-Guided Efficient Re-sampling (POGER). It worked by selecting a subset of unusual keywords, i.e., characterizing low probabilities of appearing in their contexts according to the given LLM. Then, the text is re-sampled, namely, each identified unusual keyword is removed, and the LLM is prompted to fill in the gap. If the resulting text is similar enough, the original text was likely LLM-generated.

Mitchell et al. [12] defined a new curvature-based criterion for judging if a passage is generated from a given LLM. This approach, called DetectGPT, does not require training a separate classifier or collecting a dataset of real or generated passages. It uses only log probabilities computed by the LLM

of interest and random perturbations of the passage from another generic pre-trained language model. To classify a candidate passage, DetectGPT first generates minor perturbations of the passage using a generic pre-trained model such as T5. DetectGPT compares the log probability under p of the original sample with each perturbed corresponding sample. If the average log ratio is high, the sample is likely AI-generated content. The main contribution of this work was to identify a property of the log probability function computed by a wide variety of large language models, showing that a tractable approximation to the trace of the Hessian of the model’s log probability function provides a useful signal for detecting model samples.

There are also separate approaches based on fine-tuning the LLMs to classify whether the text is AI-generated. Harrag et al. [13] fine-tuned a BERT model, specifically AraBERT, to differentiate between human-written and AI-generated Arabic tweets, primarily produced by GPT-2. They achieved very promising results with an F1-score equal to 98.7%. Rodriguez et al. [14] also trained a BERT-based model to identify texts that were fully or partially created by AI. They showed that one can fine-tune a RoBERTa model with texts from one scientific domain, and it will still accurately detect AI-generated texts from another domain, provided that a few samples from the new field are used in the fine-tuning process.

III. PROPOSED METHOD

In the current study, we followed the procedure outlined in our previous paper [4]. We used the feature-based approach and employed a binary classifier, training it to detect LLM-generated emails. As the training data, we used original emails and their LLM-generated versions. From each email, we extracted various features, such as 1) token-level perplexity (1 feature), measuring how likely the chosen LLM is to generate the input text sequence [7], 2) burstiness (1 feature), accounting for words’ distribution and occurrence patterns in a generated text [15], 3) distribution of sentence length (6 features: average, standard deviation and variance of sentence length in words and characters, 4) average word char length (1 feature), 5) punctuation metrics (2 features), counting the number of punctuation marks (.,:;!?) per number of sentences

TABLE I
TEN MOST DISCRIMINATIVE FEATURES FOR DETECTING LLM-GENERATED EMAILS, FOR ENGLISH AND POLISH LANGUAGE, SORTED BY THEIR IMPORTANCE.

Rank	English	Polish
1	number_of_errors	punctuation_per_sentence
2	no_space_after_punctuation	number_of_errors
3	stdev_sentence_char_length	stdev_sentence_char_length
4	variance_sentence_char_length	variance_sentence_char_length
5	variance_sentence_word_length	variance_sentence_word_length
6	stdev_sentence_word_length	stdev_sentence_word_length
7	double_spaces	no_space_after_punctuation
8	text_errors_by_category.typos	number_of_sentences
9	punctuation_per_sentence	text_errors_by_category.typos
10	average_word_char_length	double_spaces

and per number of characters, 6) general statistics (3 features), such as the number of characters, words, and sentences, 7) *Stylometrix* features (172 features for Polish and 196 features for English), describing stylometric characteristics obtained using the *Stylometrix* library [16]. 8) emotion-related features (5 features), such as the use of emojis, the number of question/exclamation marks, and occurrences of multiple question/exclamation marks (e.g., ??, !!,?!?).

We also extracted 26 features related to errors in text. Most of them were extracted using the Python library `language-tool-python` [17]:

- *Editing-related errors (17 features)*: features capturing a variety of typographical and stylistic errors. We counted mistakes like missing spaces after punctuation marks, double spaces between words, inconsistencies in the use of American and British English conventions, errors related to incorrect use of uppercase and lowercase letters, awkward word combinations (collocations), and incorrect word order. This category also encompasses issues such as unnecessary repetition of words, improper punctuation, and errors in forming compound words.
- *Spelling-related errors (4 features)*: we counted general spelling mistakes, probable typos, and errors involving the incorrect spelling of multi-word phrases.
- *Grammar-related errors (2 features)*: number of mistakes related to the rules of grammar, such as subject-verb agreement and sentence structure.
- *Other (3 features)*: general count of errors, miscellaneous errors, and semantic errors.

In total, 241 features for English and 217 for Polish were extracted for each email text. In this study, we aimed to assess the impact of spelling, grammar, and editing-related features on the detection of LLM-generated emails.

IV. EXPERIMENTS

In this work, we analyzed detection performance for various feature groups to find out what impact the features related to spelling, grammar, and editing have on LLMGT detection accuracy. The study setup is similar to the one used in our previous work [4]. Since our previous experiments revealed that a random forest classifier with 100 trees yielded the best results, we used it exclusively here. All experiments

were conducted using `scikit-learn` [18] library version 1.4.2, following a 10-fold cross-validation scheme.

As for the email data, we used three publicly available email datasets: “Spam email dataset” [19], containing email subjects and their content in plain text, “Email classification dataset” [20], and “The Spam Assassin Email Classification Dataset” [21]. Out of them, we obtained a set of 20156 emails.

Since we also wanted to detect email messages in Polish, we had to add our in-house data. These data contained 38776 emails, both in Polish and English. Next, we filtered out emails with less than ten characters of content and those that were created later than 2022 to be sure that none of the widely used LLMs (such as GPT-3.5) generated them. We also filtered out spam and advertisement emails to focus just on emails that can be considered as human conversations. Eventually, we obtained a dataset with 9885 original (i.e., human-written) emails in English and 471 emails in Polish.

Next, we created a “mirror” dataset with LLM-generated email texts that closely resembled the content (both in terms of the topic and the sentiment) of human-written emails. Every generated email was based on a single real email (see the generation scheme shown in Fig. 1). Through OpenAI API, we provided the email’s subject and body in plaintext and prompted `gpt-3.5-turbo-0125` model to shortly summarize the email and classify the email’s tone as either `formal`, `neutral`, `informal`. Next, we took the summary and the tone of the email, and we prompted the same model to generate a complete email based on that information. This way, we created a dataset with generated emails. Noteworthy, they were not simple paraphrases of original emails, but new emails of roughly similar size, generated based on a short summary of original emails.

V. RESULTS

We evaluated the detection ability for various feature groups using standard metrics, such as accuracy (ACC), precision, recall, F1-Score, and the area under the curve (AUC).

Table I presents the most discriminative features, identified according to the mutual information (MI) value. It confirms what was initially stated, following our previous paper [4], that the detection relied strongly on the features related to various types of mistakes: the total number of mistakes is ranked #1




Fig. 2. Histograms for `number_of_errors` feature, for a) original and b) generated emails in English.




Fig. 3. Histograms for `number_of_errors` feature, for a) original and b) generated emails in Polish.

for English and #2 for Polish. It is also very well visible in histograms displayed in Figures 2 and 3, for emails in English and Polish, respectively. They show that real (original) emails exhibit a rather uniform distribution of errors, i.e., there is a remarkable number of emails with multiple errors. In contrast to that, most of the generated emails are either error-free or exhibit a low number of mistakes. The number of errors seems to decrease exponentially.

Since users often use tools to check spelling and grammar, we wanted to see if detecting LLM-generated text on “dirty” and “clean” texts makes a remarkable difference. Table II shows the detection metrics for various groups of features used by the classifier. For English and Polish, detection accuracy and other metrics decreased when error-related text features were removed, yet the decrease was not uniform. The F1-score

for English, after removing 26 error-related features, decreased from 0.9885 to 0.9833, so the drop was only minor (around 0.5% relative). Similar minor deterioration was observed for other metrics. However, the drop for Polish was more visible: the F1-score decreased from 0.9484 to 0.9235, i.e., by 2.5% relative.

Table II displays that removing editing-related parameters (such as missing spaces, double spaces, or incorrect casing) contributed most to this drop in detection performance for Polish. As for English, all three groups had a similar, minor impact on the evaluation metrics.

We also made an interesting observation when we selected the 10 best features and trained the detection classifier in a 10-dimensional space. When using the 10 best features for English (the list is shown in Table I), we were able to create

TABLE II
RESULTS OF DETECTION OF LLM-GENERATED EMAILS FOR VARIOUS GROUPS OF FEATURES

Language	Features	# Features	Accuracy	Precision	Recall	F1 Score	ROC AUC
English	All	244	0.9882	0.9923	0.9848	0.9885	0.9995
	All but editing-related	224	0.9859	0.9901	0.9825	0.9863	0.9991
	All but spelling-related	237	0.9861	0.9905	0.9827	0.9866	0.9992
	All but grammar-related	239	0.9878	0.9916	0.9847	0.9881	0.9994
	All but error-related	218	0.9828	0.9878	0.9790	0.9833	0.9987
	10 best of all	10	0.9884	0.9931	0.9845	0.9888	0.9995
	10 best, no error-related	10	0.9700	0.9798	0.9617	0.9707	0.9948
Polish	All	220	0.9461	0.9375	0.9607	0.9484	0.9868
	All but editing-related	200	0.9192	0.8942	0.9549	0.9227	0.9773
	All but spelling-related	213	0.9389	0.9278	0.9560	0.9413	0.9857
	All but grammar-related	215	0.9410	0.9355	0.9524	0.9432	0.9876
	All but error-related	194	0.9202	0.8935	0.9575	0.9235	0.9747
	10 best of all	10	0.9202	0.9185	0.9248	0.9211	0.9774
	10 best, no error-related	10	0.8953	0.8792	0.9226	0.8989	0.9624

TABLE III
TEN MOST DISCRIMINATIVE FEATURES FOR DETECTING LLM-GENERATED EMAILS, FOR ENGLISH AND POLISH LANGUAGE, SORTED BY THEIR IMPORTANCE, ASSUMING EDITING, SPELLING, AND GRAMMAR CORRECTNESS.

Rank	Features	
	English	Polish
1	variance_word_char_length	punctuation_per_sentence
2	stdev_word_char_length	stdev_sentence_word_length
3	stdev_sentence_char_length	stdev_sentence_char_length
4	variance_sentence_char_length	variance_sentence_char_length
5	stdev_sentence_word_length	variance_sentence_word_length
6	variance_sentence_word_length	number_of_sentences
7	stylometrix_statistics_ST_SENT_D_NP	variance_word_char_length
8	stylometrix_statistics_ST_SENT_D_PP	stdev_word_char_length
9	punctuation_per_sentence	average_sentence_char_length
10	average_word_char_length	average_sentence_word_length

a detection model of the same detection efficacy as for the full 244-feature space. However, if we removed error-related features, a classifier working in such a feature space would have accuracy and the F1-score lower by 1.8% and 1.5% relative, respectively. To compensate for the loss of error-related features, at least 20-feature space would be needed (see the ACC and AUC values against the number of features shown in Fig. 4).

As for Polish, using only 10-best features would yield results lower than for the full feature space by more than 2.5%. At least 30 features would be required to achieve the accuracy results as for the full feature set (see Fig. 5). After removing error-related features, the 10 best feature space allows the detection with the metrics by around 5% relative lower than for the full set. However, using the 20 best non-error-related features seems optimal for Polish and yields better accuracy even than for the complete feature set. Yet, the results are clearly inferior to those for English, which partially can be a consequence of a much smaller size of the Polish email dataset and partially of different characteristics of the tools used for Polish.

Table III displays the 10 most discriminative non-error-related features. One can see that statistical parameters related to word and sentence length, as well as the number of punctuation marks per sentence (which is correlated with sentence length, especially for Polish), exhibited the highest discriminative power when detecting LLM-generated emails




Fig. 4. Comparison of ACC and AUC for detection with and without error-related features for English

in the absence of error-related features. As for English, also the usage statistics of noun phrases (NP) and prepositional phrases (PP) turned out to be important.

VI. CONCLUSIONS

In our paper, we have expanded upon the foundational work presented in [4], and we investigated the impact of spelling and




Fig. 5. Comparison of ACC and AUC for detection with and without error-related features for Polish

editing correctness on detection accuracy, motivated by the observed significance of these factors. Our findings reinforce the observations and provide further insights into this dependency. The code used in the experiments and the complete feature list have been made public¹.

For English texts, when using a limited number of features, we noted a slight decrease in ACC and AUC for the feature set, excluding error-related features; however, this difference became less relevant as more features were used. In contrast, for the Polish emails, the differences were more significant even when full feature lists were used, indicating a decline in prediction quality by around 2% relative for ACC, AUC, and F1-scores. The results indicate that the proposed detection method can be equally effective for English even if spelling- and grammar-checking tools are used. As for Polish, to compensate for the “loss” of text errors, we need to use more features and, potentially, also to seek new ones.

REFERENCES

- [1] Europol, “Chatgpt: the impact of large language models on law enforcement,” 2023. doi: 10.2813/255453
- [2] H. Williams and C. McCulloch, “Truth decay and national security: Intersections, insights, and questions for future research,” Santa Monica, CA, USA, 2023. [Online]. Available: <https://www.rand.org/pubs/perspectives/PEA112-2.html>
- [3] K. T. Gradoń, “Generative artificial intelligence and medical disinformation,” *British Medical Journal*, no. 384, 2024. doi: 10.1136/bmj.q579
- [4] P. Gryka, K. Gradoń, M. Kozłowski, M. Kutyla, and A. Janicki, “Detection of AI-generated emails – a case study,” in *Proc. 13th International Workshop on Cyber Crime (IWCC 2024)*, Vienna, Austria, 2024, (accepted for publication).
- [5] A. Knott, D. Pedreschi, R. Chatila, T. Chakraborti, S. Leavy, R. Baeza-Yates, D. Eysers, A. Trotman, P. D. Teal, P. Biecek, S. Russell, and Y. Bengio, “Generative AI models should include detection mechanisms as a condition for public release,” *Ethics and Information Technology*, vol. 25, no. 4, p. 55, 12 2023. doi: 10.1007/s10676-023-09728-4
- [6] K. Krishna, Y. Song, M. Karpinska, J. Wieting, and M. Iyer, “Paraphrasing evades detectors of AI-generated text, but retrieval is an effective defense,” *Advances in Neural Information Processing Systems*, vol. 36, 3 2024.
- [7] F. Jelinek, R. L. Mercer, L. R. Bahl, and J. K. Baker, “Perplexity—a measure of the difficulty of speech recognition tasks,” *The Journal of the Acoustical Society of America*, vol. 62, no. S1, pp. S63–S63, 1977. doi: 10.1121/1.2016299
- [8] M. Chakraborty, S. T. I. Tonmoy, S. M. M. Zaman, S. Gautam, T. Kumar, K. Sharma, N. Barman, C. Gupta, V. Jain, A. Chadha, A. Sheth, and A. Das, “Counter Turing test (CT2): AI-generated text detection is not as easy as you may think - introducing AI detectability index (ADI),” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023. doi: 10.18653/v1/2023.emnlp-main.136 pp. 2206–2239. [Online]. Available: <https://aclanthology.org/2023.emnlp-main.136>
- [9] I. Cingilloglu, “Detecting AI-generated essays: the ChatGPT challenge,” *International Journal of Information and Learning Technology*, vol. 40, pp. 259–268, 5 2023. doi: 10.1108/IJILT-03-2023-0043
- [10] L. Fröhling and A. Zubiaga, “Feature-based detection of automated language models: tackling GPT-2, GPT-3 and Grover,” *PeerJ Computer Science*, vol. 7, p. e443, 4 2021. doi: 10.7717/peerj.cs.443
- [11] Y. Shi, Q. Sheng, J. Cao, H. Mi, B. Hu, and D. Wang, “Ten words only still help: Improving black-box AI-generated text detection via proxy-guided efficient re-sampling,” *arXiv preprint*, vol. arXiv:2402.09199, 2024. [Online]. Available: <http://arxiv.org/abs/2402.09199>
- [12] E. Mitchell, Y. Lee, A. Khazatsky, C. D. Manning, and C. Finn, “DetectGPT: Zero-shot machine-generated text detection using probability curvature,” in *Proc. International Conference on Machine Learning*. Online: PMLR, 2023, pp. 24 950–24 962.
- [13] F. Harrag, M. Dabbah, K. Darwish, and A. Abdelali, “Bert transformer model for detecting Arabic GPT2 auto-generated tweets,” in *Proceedings of the Fifth Arabic Natural Language Processing Workshop*, I. Zitouni, M. Abdul-Mageed, H. Bouamor, F. Bougares, M. El-Haj, N. Tomeh, and W. Zaghouani, Eds. Barcelona, Spain (Online): Association for Computational Linguistics, Dec. 2020, pp. 207–214. [Online]. Available: <https://aclanthology.org/2020.wanlp-1.19>
- [14] J. D. Rodriguez, T. Hay, D. Gros, Z. Shamsi, and R. Srinivasan, “Cross-domain detection of GPT-2-generated technical text,” in *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, vol. 2022.naacl-main.88. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.naacl-main.88 pp. 1213–1233.
- [15] S. Mukherjee, “Exploring burstiness: Evaluating language dynamics in LLM-generated texts,” 2023, [Online]. Available: <https://ramblersm.medium.com/exploring-burstiness-evaluating-language-dynamics-in-llm-generated-texts-8439204c75c1> (Accessed on Apr 30, 2024).
- [16] I. Okulska, D. Stetsenko, A. Kołos, A. Karlińska, K. Głabińska, and A. Nowakowski, “Stylometrix: An open-source multilingual tool for representing stylometric vectors,” *arXiv preprint arXiv:2309.12810*, vol. 2309.12810, 9 2023.
- [17] J. Morris, “LanguageTool Python library,” 2024, <https://pypi.org/project/language-tool-python/> (Accessed on May 10, 2024). [Online]. Available: <https://pypi.org/project/language-tool-python/>
- [18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [19] _w1998, “Spam email dataset,” 2023, (Accessed on Jan 14, 2024). [Online]. Available: <https://www.kaggle.com/datasets/jacksonscie/spam-email-dataset/data>
- [20] R. Modi, “Email classification dataset,” 2023, (Accessed on Jan 14, 2024). [Online]. Available: <https://github.com/rmodi6/Email-Classification/tree/master>
- [21] Apache Public Datasets, “The Spam Assassin Email Classification Dataset,” 2023, (Accessed on Jan 14, 2024). [Online]. Available: <https://www.kaggle.com/datasets/ganiyuolalekan/spam-assassin-email-classification-dataset/data>

¹<https://github.com/mksochota16/anti-gpt-checker/>

A network clustering method based on intersection of random spanning trees

László Hajdu^{*†}, András London[‡], and András Pluhár[‡]

^{*}Innorenue CoE, Livade 6a, SI-6310 Izola, Slovenia

[†]University of Primorska, Faculty of Mathematics, Natural Sciences and Information Technologies

Glagoljaška 8, SI-6000 Koper, Slovenia, <https://orcid.org/0000-0002-1832-6944>

[‡]University of Szeged, Institute of Informatics, H-6726 Árpád tér 2, Hungary

Abstract—We use a special edge centrality measure for node clustering in complex networks. The measure is based on the ‘spanning tree intersection’ value motivated by previous work on the intersection and minimum expected overlap of random spanning trees in complex networks. First, we show that this new metric differs from some well-known edge centralities on random network models and real-world networks. Then, we show the applicability of the metric for clustering the nodes and point out some advantages over some other edge centrality based hierarchical clustering methods.

I. INTRODUCTION

IN NETWORK science, paths and shortest paths between network nodes are of extraordinary interest due to the enormous number of straightforward applications [1], [2]. Besides the applications and algorithm development [3], [4], several metrics have been derived from paths and shortest paths, such as betweenness and closeness centrality, k-path centrality, etc., and used for different purposes. For example, some community detection algorithms have efficiently used path-based edge centrality measures, e.g., [5], [6], [7].

A widely used centrality based on shortest paths is *edge betweenness*, which measures the importance of an edge in a network by the number of shortest paths that pass through the edge. A shortest path between two nodes in an unweighted network is a path that has the minimum number of edges. The edge betweenness of an edge is then calculated as the fraction of the shortest paths between all pairs of nodes in the network that pass through that edge. The importance of edges with high edge betweenness is obvious. They are vital in connecting different network parts and facilitating communication and information flow.

On the other hand, as many real-world situations show, there is “something artificial about shortest paths. It seems that shortest paths are sometimes too short and do not correspond to the underlying natural logic of the network”, see [1] for more details. Information or traffic does not always prefer the shortest paths, but it sometimes also takes much longer paths. Furthermore, betweenness centrality, a widely used metric that is based on shortest paths, is unstable: adding just one ‘shortcut’ link to the network can dramatically change the scores of edges in the network [8].

Such considerations have led to an alternative model of edge importance based on randomly selected spanning trees [9], [10]. A *spanning tree* of a connected network of n nodes is a tree composed of $n - 1$ edges such that they connect all n nodes in the network. Many results support the intuition of using spanning trees to compute edge centrality. In social networks, the spread of information follows tree-like cascades [11]; in technological (e.g., roads, electricity) and transaction networks (e.g., financial transactions, trade), the routes used often show tree-like structures [12].

Spanning Centrality is a measure of the importance of an edge in a network based on the number of spanning trees that contain the edge. The spanning centrality of an edge is calculated as the proportion of all spanning trees that contain the edge. Edges with high spanning centrality are essential because they are involved in many different paths through the network. This means they play a vital role in connecting other parts of the network and facilitating communication and information flow. Spanning centrality can be used to analyze the structure of networks and identify bottlenecks or critical points that could disrupt the flow of information or communication [13].

Recently it has been discussed that the average intersection (i.e., the number of common edges) of random spanning trees encodes some structural information about the network regarding homogeneity [14] and resilience [15], and also closely related to several concepts like *fairest edge usage*, *spanning tree modulus* and *secure broadcasting* over networks, see [14], [16]. The idea of checking that an edge of a network will likely be in the intersection of randomly chosen spanning trees then comes naturally.

In this short paper, we introduce the concept of *spanning tree intersection edge centrality* with the goal of using it in top-down hierarchical clustering of network nodes. This is motivated by at least two things. One is to use a metric, instead of betweenness, which is faster to calculate, but possibly provide a similar clustering. The other is to see how much it gives different clusters than the usual clustering or community detection procedures.

Throughout this paper, $G = (V, E)$ will be a finite, connected, undirected, and unweighted graph representing a network with $|V| = n$ nodes and $|E| = m$ edges.




Fig. 1. Running time comparison on Gaussian random partition network.




Fig. 2. Running time comparison on Watts-Strogatz network.

II. DEFINITIONS AND METHODS

A. Edge centrality measures

Given a network G , edge centrality is a function that assigns a number to each edge of the network. Next, we define the three edge centrality measures we want to compare and briefly discuss the algorithms' complexity for calculating them.

Edge betweenness: Let σ_{ij} be the number of shortest paths between nodes i and j of an undirected and connected graph $G = (V, E)$, and $\sigma_{ij}(e)$ be the number of shortest paths between them that pass through edge $e \in E$. Note that there can be more than one shortest path between any two nodes in the case of unweighted graphs. The betweenness centrality $BC(e) \in (0, 1]$ of edge e is defined as $BC(e) = \sum_{i \neq j} \sigma_{ij}(e) / \sigma_{ij}$. The greater the number of shortest paths that pass through a particular edge, the greater the importance of that edge. Intuitively, a high $BC(e)$ identifies edges that act as bridges or intermediaries between nodes in the network and are essential for maintaining efficient communication and information flow between different parts of the network.

Spanning tree centrality: Given an undirected and connected graph $G = (V, E)$, the spanning centrality $SC(e) \in (0, 1]$ of an edge $e \in E$ is defined as the fraction of spanning trees of G that contain e . Intuitively, a high $SC(e)$ quantifies how important the edge e is for G to ensure its connectivity. An edge with a high SC value is present in most spanning trees, all of which will fall apart if the edge is removed from G . In the case where $SC(e) = 1$, G is broken when e is removed, i.e., e is a cutting edge if and only if $SC(e) = 1$. This means that such an edge participates in all possible spanning trees. Important edges participate in many spanning trees, assuming that spanning trees encode candidate paths through which information flows.

Spanning tree intersection centrality: The *spanning tree intersection centrality* value $STI(e) \in (0, 1]$ of an edge $e \in E$ is the fraction of spanning tree pairs among all such pairs that both contain e , i.e., e is in the intersection of the pair.

In other words, $STI(e)$ shows how likely an edge will be in the intersection of randomly chosen spanning tree pairs. Intuitively, STI is very similar to SC but more restrictive in giving a high value for an edge. It also holds here that $STI(e) = 1$ if and only if e is a cutting edge. The relevancy of edges with high STI values can be seen from the perspective of spanning tree overlap. These are the edges that are somehow crucial in order to get a not empty overlap between randomly chosen spanning trees.

According to [15], the minimum of the expected number of edges in the intersection of two randomly chosen spanning trees is $(n-1)^2/m$, while the expected value can be calculated precisely as $\sum_{e \in E} p^2(e)$, where $p(e)$ is the probability that the edge e is in a uniform random spanning tree. It suggests (not dealing with the variance) that the probability of an edge e being in the intersection of randomly chosen spanning trees is $\approx p^2(e)$. In this sense, STI can be derived from SC , but we argue that it can provide some added value to it, and more restrictive in assigning high values to the edges.

B. Algorithms

Betweenness centrality can be computed efficiently, e.g., using the Brandes algorithm, in $O(nm)$ time.

The number of spanning trees of G is explicitly known by Kirchhoff's matrix tree theorem [17] and can be computed as the product of the positive eigenvalues of the graph Laplacian L divided by n . It can be calculated in polynomial, $(O(n^3))$, time (let \tilde{L} be the Laplacian with the Laplacian with row and column i^{th} removed (for any i), then $\det \tilde{L}$ be the number of spanning trees). This allows us to precisely define *uniform random spanning trees*, i.e., the uniform probability distribution over all spanning trees.

For an edge $e \in E(G)$, we can compute the probability $p(e)$, that the edge e is in a uniform random spanning tree of G , in polynomial time. This is $p(e) = \det(\tilde{L}_{G \setminus \{e}}) / \det \tilde{L}$, where $\tilde{L}_{G \setminus \{e}}$ is defined similarly as above for the graph $G \setminus e$, the

graph obtained by contracting the original one after the edge e is deleted. So, the nominator is the number of spanning trees containing e .

Although we can determine the probability of picking a uniform spanning tree, it does not follow that we can generate one. Moreover, calculating determinants is still expensive. Several algorithms have been developed to generate random spanning trees of an undirected graph; see, for instance, Broder's [18] and Wilson's [19] algorithms, worst case $O(nm)$, but on average $O(n \log n)$ time. Moreover, almost linear time algorithms exist [20]. That allows us to perform Monte-Carlo-style experiments, described in detail next.

Figure 1 shows the running times of the betweenness centrality and our spanning tree intersection centrality in a Gaussian random partition network. The network was generated with the size parameter n varying from 100 to 4000 in increments of 100. Each network was constructed with sub-community sizes $s = n/10$ and $v = n/20$, and probabilities of intra-community and inter-community connections set at 0.7 and 0.001, respectively. We also generated Watts-Strogatz networks where the size parameter varied from 100 to 4000 in increments of 100. For each network, we set the number of nearest neighbours in the ring topology (k) to 4 and the rewiring probability (p) to 0.1. The corresponding running times can be seen on Figure 2.

The running times for many other random network models have a similar shape. Our choice of the two models we present was based on the motivation of using the metric for community detection purposes, and these two models provide well-structured networks with community structure and other small-world properties that appear in real-world networks.

III. EXPERIMENTAL RESULTS

A. Experiment design

The main objective of the simulation was to explore the behavior of the two spanning tree-based centrality measures and compare them with the well-known and widely used edge betweenness centrality. As we noted the direct connection between SC and STI before, we restrict the presentation of the comparison results to those between BC and STI . To perform a detailed comparison, we collected real-world networks with varying properties and a curated set of random model networks. In the following section, we present the selection of networks we used in the simulation.

Selection of random and real-world networks: During the initialization of the experiment, we generated four different types of networks with the following parameters:

- 1) Random regular graph with parameter $d = 3, 4, \dots, 9$ (degree of each node) and number of nodes $n = 100, 200, 500$ [21], [22].
- 2) Erdős-Rényi random graph with parameter $p = 0.1, 0.2, \dots, 1$ (probability of edge creation) and number of nodes $n = 100, 200, 500$ [23], [24].
- 3) Preferential attachment random network using the `barabasi_albert_graph` function from NetworkX with parameter $m = 1, 2, \dots, 10$ (number of edges to attach

Algorithm 1 Random Spanning Tree Simulation (RSTS)

```

1: Input: Graph  $G(V, E)$ , sample size  $s$ 
2:  $j \leftarrow 0$ 
3:  $\forall e \in E : C_e^{\text{intersection}} = 0$ 
4: While  $j < s$ 
5:    $H1 = \text{WilsonRST}(G)$ 
6:    $H2 = \text{WilsonRST}(G)$ 
7:    $I = \text{intersection}(H1, H2)$ 
8:   For  $e$  in  $E$ 
9:     If  $e \in I : C_e^{\text{intersection}} \leftarrow C_e^{\text{intersection}} + 1$ 
10:  End For
11: End While
12:  $\forall e \in E : C_e^{\text{intersection}} \leftarrow C_e^{\text{intersection}} / s$ 

```

from a newly created node to existing nodes) and number of nodes $n = 100, 200, 500$ [25].

- 4) Watts-Strogatz network using the `watts_strogatz_graph` function from NetworkX with parameters $k = 4$ (connected nearest neighbors in the ring topology) and $p = 0.1, 0.2, \dots, 1$ (probability of rewiring the edges) and number of nodes $n = 100, 200, 500$ [26].

For each type and parameter, we generated 100 networks and then calculated sample means and standard deviations.

In the case of real-world networks, we selected 15 well-known networks from the field of network science with various graph properties. The list of the selected graphs and their important basic properties can be seen in Table 1.

Tools and Libraries: The simulation environment was implemented in Python 3.9 using the following libraries. NetworkX (2.6.3) library was used for the betweenness centrality calculation, random network generation, and other network manipulation. Our own modified version of the open-source library DPPy (0.3.2) [27] was used to extract the spanning trees. However, the Wilson algorithm for the random spanning tree generation (sample method) remained unchanged. (The original library could not return the network graph containing the random spanning tree itself.). We visualized the results using the Matplotlib (3.5.2) and Seaborn (0.12.1) libraries. Pandas (1.4.2) was used for basic data manipulation, while Numpy (1.22.3) was used for basic calculations regarding the results.

Simulation Environment: The pseudocode shows the steps of our Random Spanning Tree Simulation environment. In the case of a network where we had multiple connected components, we always chose the largest connected component.

We generated a $2s$ number of random spanning trees during the simulation. The algorithm counts how often it is in the intersection of two random spanning trees. To calculate the expected value of being the intersection, we divide the result by the sample size for each edge. Throughout our experiment, we used $s = 1000$ as the sample size parameter in the case of both random and real networks (except for the real Email-Enron network, due to its large size, where we used $s = 100$).




Fig. 3. Correlations between Intersection Probability and Edge Betweenness measures.

B. Comparison on random network models

As mentioned above, we generated 100 networks for each type and parameter. To gain valuable insights from the simulation, we took the mean $STI(e)$ and $BC(e)$ values (i.e., “intersection probabilities” and edge betweenness) for each edge, computed their correlations for each graph, and then aggregated them across the different parameters, sizes, and graph types to compare the patterns and behavior of the different measures.

The heatmap in Figure 3 shows the correlation between intersection probability and edge betweenness. In this case, in the case of Erdős-Rényi networks, no significant pattern can be observed. In the case of Barabási-Albert networks, there is a negative correlation with increasing size and k parameter values. The values are between -0.01 and -0.58. It can be seen that in the case of size 500, the values are between -0.39 and -0.55, so there is a negative correlation even at higher k-parameters. Random regular and Watts-Strogats networks again behave similarly for small d and p parameters. However, at higher parameter values, the correlation becomes negative in the case of WS. On the other hand, the random regular shows a slight decrease in correlations with decreasing size and parameter, with values between 0.05 and 0.59.

C. Comparison on real-world networks

This section presents, analyzes, and interprets our results for the real-world networks shown in Table 1 and introduces the behavior of the different measures on an arbitrary real network (as Figure 4 shows). The table briefly overviews various network characteristics and metrics for the given set of real-world networks. These metrics include the number of nodes, edges, network density, clustering coefficient, average path length, expected intersection (equals to $(n - 1)^2/m$),




Fig. 4. Comparison intersection probability measure with the edge betweenness on the Florentine graph. (left: Intersection Probability, right: Edge Betweenness) Thicker edge refers to higher edge centrality.

observed intersection (observed intersection value based on the simulations), std (standard deviation), the adjusted index

$$RTI = \frac{\text{observed mean} - \text{min.expected}}{n - \text{min.expected}},$$

modularity, and the same correlation pairs as in the case of the random network in the previous section. More precise definitions of the calculated properties can be found in [28].

As an example shown in Figure 4, the thickness of the edge shows the actual value of the given measure; more precisely, the thicker the edge between two nodes, the higher the corresponding measure of the edge.

Table 1. shows the results on each of the real networks we used during the evaluation. The results show that our intersection probability and spanning tree probability measures have a relatively higher correlation with edge betweenness in the case of “Florentine”, “Adjnoun”, “Jazz” and “C-elegans” networks. As the previous section shows, our two measures are also highly correlated with each other in real networks, which means they express similar properties of the edges in the network.

The relationship between the correlation between spanning tree intersection centrality and betweenness centrality with Newman modularity (left) and RTI (right), respectively, on the investigated real-world dataset is shown in Figure 5. It suggests that the modularity and RTI are higher if the correlation between the two different centralities is high and lower in the case of a lower correlation. It would be worth investigating this effect in the future, as it may help us better understand the mesoscale structure of networks (more details will be given below).

IV. COMMUNITY DETECTION BASED ON EDGE IMPORTANCE

A well-known community detection algorithm proposed by Newman and Girvan [6] uses centrality indices to find community boundaries. It is assumed that communities or groups are only loosely connected by a few edges between groups. Therefore, all shortest paths between communities must be along one of these few edges. Then, the edges connecting communities should have a high edge betweenness.




Fig. 5. Scatter plot that shows the connection of Intersection Probability - Edge betweenness correlation with Modularity measure and *RTI*.

By iteratively removing these edges, the communities can be separated at different hierarchical levels.

A natural idea is to change edge betweenness to spanning centrality, spanning tree intersection centrality, or a metric derived from these. This should be done carefully, as dangling edges (connected to 1 degree nodes) always have *STI* (and *SC*) values equal to one (see Figure 4). To handle such nodes, one can perform a preprocessing step that removes these nodes after automatically assigning the label of their only neighbor (which label will be assigned later in the process), or a unique label, depending on the specified requirement for the clustering algorithm. After this step, the edges are removed one by one according to the decreasing order of *STI* values. At the end of the process, all edges are removed. If we think of this process as a top-down hierarchical procedure, then the cutters are defined as the connected components of the network at a particular hierarchical level. The pseudocode of the algorithm is given below. Fig. 4 shows a benchmark example with the clusters found by the algorithm. Since the cluster structure is determined at a particular hierarchical level (i.e. after removing a certain number of edges to get separate components of the network), the level should be specified. This can be done, for instance maximizing the Newman modularity function [29],

Algorithm 2 Iterative edge removal based on *RTI* values

```

1: Input: Graph  $G(V, E)$ ,
2: Repeat
3:   For  $e$  in  $E$ 
4:     If  $RTI(e) = 1$ ,  $e = (u, v)$  is a dangling edge with  $d(v) = 1$  then  $C(v) = C(u)$ , remove  $v$ 
5:     End If
6:     If  $RTI(e) > \max(RTI)$  then
7:        $\max(RTI) = RTI(e)$ 
8:        $\text{argmax}(STI) = e$ 
9:     End If Remove  $e$  from  $G$ 
10:  End For
11: Until there is no more edge in  $G$ 
12: Output: cluster dendrogram
    
```

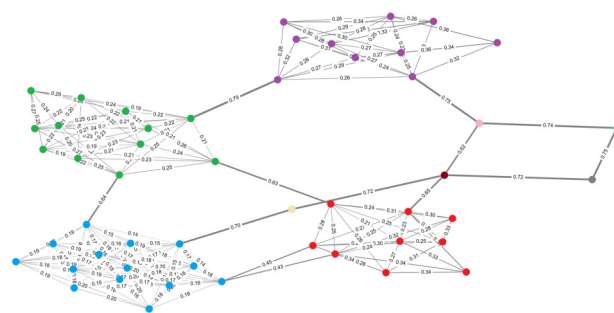


Fig. 6. Clustering given by the *RTI* value based top-down hierarchical clustering method. The network was generated by the gaussian_random_partition_graph() method from the NetworkX package.

that can be calculated for any given clustering.

V. SUMMARY

This paper introduced an edge centrality metric based on random spanning tree intersections. We presented an algorithm to compute it and compared it with the widely used edge betweenness and spanning centrality metrics. Our initial results suggest that this metric may be helpful in determining essential links of the network in terms of path usage, connectivity and resilience. We also experimented that this metric is efficiently applicable for clustering the nodes and have some advantages over some other edge centrality based top-down hierarchical clustering methods. We also hypothesize that a metric derived from betweenness and spanning tree intersection centrality could help to optimize the modularity in the Girvan-Newman algorithm. The exploration of this can be the topic of a future work.

Acknowledgement. András London was supported by National Research, Development and Innovation Office—NKFIH Fund No. SNN-135643. László Hajdu acknowledges the European Commission for funding the InnoRenew CoE project (Grant Agreement no. 739574) under the Horizon2020 Widespread-Teaming program and the Republic of Slovenia (Investment funding of the Republic of Slovenia and the European Union of the European Regional Development Fund). He is also grateful for the support of the

Slovenian Research Agency (ARRS) through grant N2-0171 and for the support of University of Primorska through postdoc grant No. 2991-10/2022.

REFERENCES

- [1] A. Gulyás, Z. Heszberger, and J. Biró, *Paths: Why is life filled with so many detours?* Springer Nature, 2021, doi:10.1007/978-3-030-47545-1.
- [2] A. Gulyás, Z. Heszberger, J. Biró, A. Gulyás, Z. Heszberger, and J. Biró, “The universal nature of paths,” *Paths*, pp. 45–65, 2021.
- [3] A. Madkour, W. G. Aref, F. U. Rehman, M. A. Rahman, and S. Basalamah, “A survey of shortest-path algorithms,” *arXiv preprint arXiv:1705.02044*, 2017.
- [4] L. M. Laskov and M. L. Marinov, “List of pareto optimal solutions of a biobjective shortest path problem,” in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 603–613, doi:10.15439/2023F3718.
- [5] A. Bóta, M. Krész, and A. Pluhár, “Dynamic communities and their detection,” *Acta Cybernetica*, vol. 20, pp. 35–52, 2011.
- [6] M. Girvan and M. E. Newman, “Community structure in social and biological networks,” *PNAS*, vol. 99, no. 12, pp. 7821–7826, 2002, doi:10.1073/pnas.122653799.
- [7] P. G. Sun and Y. Yang, “Methods to find community based on edge centrality,” *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 9, pp. 1977–1988, 2013, doi:10.1016/j.physa.2012.12.024.
- [8] M. E. Newman, “A measure of betweenness centrality based on random walks,” *Social Networks*, vol. 27, no. 1, pp. 39–54, 2005, doi:10.1016/j.socnet.2004.11.009.
- [9] X. Qi, E. Fuller, R. Luo, and C.-q. Zhang, “A novel centrality method for weighted networks based on the kirchhoff polynomial,” *Pattern Recognition Letters*, vol. 58, pp. 51–60, 2015, doi:10.1016/j.patrec.2015.02.007.
- [10] A. S. Teixeira, P. T. Monteiro, J. A. Carriço, M. Ramirez, and A. P. Francisco, “Spanning edge betweenness,” in *Workshop on Mining and Learning with Graphs*, vol. 24, 2013, pp. 27–31.
- [11] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, “Inferring networks of diffusion and influence,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 5, no. 4, pp. 1–37, 2012.
- [12] S. Lämmer, B. Gehlsen, and D. Helbing, “Scaling laws in the spatial structure of urban road networks,” *Physica A: Statistical Mechanics and its Applications*, vol. 363, no. 1, pp. 89–95, 2006, doi:10.1016/j.physa.2006.01.051.
- [13] C. Mavroforakis, R. Garcia-Lebron, I. Koutis, and E. Terzi, “Spanning edge centrality: Large-scale computation and applications,” in *Proceedings of the 24th International Conference on World Wide Web*, 2015, pp. 732–742, doi:10.1145/2736277.2741125.
- [14] N. Albin, J. Clemens, D. Hoare, P. Poggi-Corradini, B. Sit, and S. Tymochko, “Fairest edge usage and minimum expected overlap for random spanning trees,” *Discrete Mathematics*, vol. 344, no. 5, p. 112282, 2021.
- [15] A. London and A. Pluhár, “Intersection of random spanning trees in small-world networks,” in *International Conference on Complex Networks and Their Applications*. Springer, 2022, pp. 337–345, doi:10.1007/978-3-031-21131-7_26.
- [16] K. Kottegoda, *Spanning tree modulus and secure broadcast games*. PhD dissertation, Kansas State University, 2020.
- [17] E. W. Weisstein, “Matrix tree theorem,” <https://mathworld.wolfram.com/>, 2000.
- [18] A. Z. Broder, “Generating random spanning trees,” in *FOCS*, vol. 89, 1989, pp. 442–447, doi:10.1109/SFCS.1989.63516.
- [19] D. B. Wilson, “Generating random spanning trees more quickly than the cover time,” in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 296–303, doi:10.1145/237814.237880.
- [20] A. Schild, “An almost-linear time algorithm for uniform random spanning tree generation,” in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 2018, pp. 214–227.
- [21] J. H. Kim and V. H. Vu, “Generating random regular graphs,” in *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, ser. STOC '03, New York, NY, USA, 2003, p. 213–222.
- [22] A. Steger and N. C. Wormald, “Generating random regular graphs quickly,” *Combinatorics, Probability and Computing*, vol. 8, no. 4, pp. 377–396, 1999, doi:10.1017/S0963548399003867.
- [23] P. Erdős and A. Rényi, “On random graphs i,” *Publicationes Mathematicae Debrecen*, vol. 6, pp. 290–297, 1959, doi:10.1515/9781400841356.38.
- [24] E. N. Gilbert, “Random graphs,” *The Annals of Mathematical Statistics*, vol. 30, no. 4, pp. 1141–1144, 1959.
- [25] A.-L. Barabasi and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, pp. 509–512, 1999, doi:10.1126/science.286.5439.509.
- [26] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, pp. 440–442, 1998, doi:10.1038/30918.
- [27] G. Gautier, G. Polito, R. Bardenet, and M. Valko, “DPPy: DPP Sampling with Python,” *Journal of Machine Learning Research - Machine Learning Open Source Software (JMLR-MLOSS)*, 2019, <http://github.com/guilgautier/DPPy/>.
- [28] A. London and A. Pluhár, “Intersection of random spanning trees in complex networks,” *Applied Network Science*, vol. 8, no. 1, p. 72, 2023.
- [29] M. E. Newman, “Modularity and community structure in networks,” *PNAS*, vol. 103, no. 23, pp. 8577–8582, 2006.

TABLE I
NETWORK STATISTICS AND RESULTS FOR SOME REAL-LIFE NETWORKS

name	nodes	edges	density	clustering	avg path length	expected	observed	std	RTI	modularity	IntersectP		SptreeP	
											EdgeB	EdgeB	EdgeB	EdgeB
Florentine	15	20	0.1905	0.1915	2.4857	9.80	10.527	1.007	0.140	0.3987	0.4945	0.5278	0.5278	0.5278
Zachary	34	78	0.1390	0.2557	2.4082	13.96	15.313	2.368	0.067	0.3807	0.0283	-0.0135	-0.0135	-0.0135
Dolphins	62	159	0.0841	0.3088	3.3570	23.40	28.980	2.847	0.145	0.4955	0.2093	0.2414	0.2414	0.2414
Iceland	75	114	0.0411	0.1573	3.1996	48.04	57.191	2.351	0.340	0.5742	0.0470	0.0099	0.0099	0.0099
Adajoun	112	425	0.0684	0.1569	2.5556	28.99	40.050	3.970	0.133	0.2947	0.4998	0.4174	0.4174	0.4174
Jazz	198	2742	0.1406	0.5203	2.2350	14.15	24.759	3.926	0.058	0.4389	0.4355	0.4458	0.4458	0.4458
C-elegans	297	2148	0.0489	0.1807	2.4553	40.79	59.926	5.598	0.075	0.3692	0.3979	0.2964	0.2964	0.2964
NetSci	379	914	0.0128	0.4306	6.0419	156.33	184.004	7.553	0.124	0.8383	0.0287	0.0298	0.0298	0.0298
Wiki-Vote	889	2914	0.0074	0.1273	4.0962	270.61	441.456	9.888	0.276	0.5470	0.1165	0.0726	0.0726	0.0726
EuroRoad	1039	1305	0.0024	0.0353	18.3951	825.63	857.468	6.501	0.149	0.8640	-0.1318	-0.1317	-0.1317	-0.1317
Polblogs	1222	16717	0.0224	0.2260	2.7375	89.18	310.507	10.450	0.195	0.4269	0.6050	0.6123	0.6123	0.6123
Arxiv GR-QC	4158	13428	0.0016	0.6289	6.0494	1286.91	2139.144	22.535	0.297	0.7886	0.0592	0.0784	0.0784	0.0784
Erdős	4991	7428	0.0006	0.0420	5.5116	3352.19	4187.384	16.113	0.510	0.7568	-0.3451	-0.3903	-0.3903	-0.3903
Power grid	4941	6594	0.0005	0.1032	18.9892	3700.88	3940.939	16.066	0.194	0.9346	-0.1219	-0.1167	-0.1167	-0.1167
Email-Enron	33696	180811	0.0003	0.0851	4.0252	6279.23	16845.420	52.275	0.385	0.5092	0.1210	0.1029	0.1029	0.1029

Efficient Maritime Healthcare Resource Allocation Using Reinforcement Learning

Tehreem Hasan
Quaid-i-Azam University
Islamabad Pakistan
Email: tehreemhasan@ele.qau.edu.pk

Farwa Batool
Quaid-i-Azam University
Islamabad Pakistan
Email: farwabatool@ele.qau.edu.pk

Mario Fiorino
Politecnico di Torino, Italy
Email: mario.fiorino@polito.it
ORCID: 0009-0007-9393-7095

Giancarlo Tretola
Department of Computer Engineering
Università Giustino Fortunato
Benevento Italy
Email: g.tretola@unifortunato.eu

Musarat Abbas
Quaid-i-Azam University
Islamabad Pakistan
Email: mabbas@qau.edu.pk

Abstract—The allocation of healthcare resources on ships is crucial for safety and well-being due to limited access to external aid. Proficient medical staff on board provide a mobile healthcare facility, offering a range of services from first aid to complex procedures. This paper presents a system model utilizing Reinforcement Learning (RL) to optimize doctor-patient assignments and resource allocation in maritime settings. The RL approach focuses on dynamic, sequential decision-making, employing Q-learning to adapt to changing conditions and maximize cumulative rewards. Our experimental setup involves a simulated healthcare environment with variable patient conditions and doctor availability, operating within a 24-hour cycle. The Q-learning algorithm iteratively learns optimal strategies to enhance resource utilization and patient outcomes, prioritizing emergency cases while balancing the availability of medical staff. The results highlight the potential of RL in improving healthcare delivery on ships, demonstrating the system’s effectiveness in dynamic, time-constrained scenarios and contributing to overall maritime safety and operational resilience.

I. INTRODUCTION

THE allocation of healthcare resources is an important and critical task for provision of quality health services [1]. This task in the restricted and often isolated setting of ship is not simply a matter of convenience; rather, it is an essential requirement that directly influences the safety and well-being of all individuals on board[2]. Unlike on land, where medical facilities are usually easily accessible, ships operate in environments where the availability of external aid can be significantly limited or delayed. Consequently, the distribution of healthcare resources becomes not just significant but paramount for mitigating risks and ensuring the uninterrupted continuation of maritime activities.

In maritime settings, the presence of proficient medical staff on board is comparable to having a mobile healthcare facility[3]. Physicians, nurses, and paramedics have crucial roles, accountable not only for immediate treatment during crises but also for preserving overall health and wellness throughout journeys. Their expertise, coupled with a variety of medical services ranging from basic first aid to complex

procedures, establishes the foundation of a ship’s healthcare framework[4].

Maritime healthcare encounters challenges beyond the provision of services. Efficient resource allocation requires ongoing monitoring of the deployment of medical personnel[5]. This includes ensuring sufficient staff numbers strategically located to promptly respond to emergencies anywhere on the ship. It also involves forecasting changes in demand based on the duration of the voyage, the nature of the cargo, and the demographics of the crew and passengers[6].

Furthermore, the distribution of healthcare resources surpasses mere logistics; it entails incentivizing effective decision-making. Immediate incentives and delayed penalties act as stimulants for proactive resource management, cultivating a culture of safety and accountability[7]. By acknowledging and reinforcing positive actions, ship operators ensure effective resource utilization, enhancing the overall resilience of the healthcare system[8].

Fundamentally, the allocation of healthcare resources in maritime settings demands careful planning, constant monitoring, and proactive decision-making[9]. It demonstrates the flexibility and resourcefulness of maritime experts navigating intricate connections among personnel, services, data, and measures to safeguard health and well-being.

Moreover, onboard medical facilities serve broader objectives of safety, security, and operational effectiveness[10]. They function as crucial support systems during crises, mitigating the impacts of adverse events. Nevertheless, the distinctive maritime environment presents challenges such as limited space, harsh weather conditions, and isolation, which magnify medical risks. Therefore, resource allocation must address these challenges, ensuring the preparedness of personnel to deliver efficient care[11].

Additionally, the distribution of healthcare resources on ships intertwines with risk management and compliance with regulations[12]. Maritime authorities impose stringent criteria on medical care provision and facility upkeep. The failure to

meet requirements may result in severe outcomes, compelling operators to comply with regulatory standards and industry best practices, thereby safeguarding the health and safety of individuals aboard.

The paper organized as follows: **Section 2** provides the Related Works. In **Section 3**, the problem definition and background of the healthcare system on ships are described. The system model and experimental setup, including the Q-learning algorithm, are presented in **Section 4**. Computational results based on simulated scenarios are provided in **Section 5**. Finally, **Section 6** concludes the paper.

II. RELATED WORKS

This section presents some state of the art on the use of AI-empowered solutions for healthcare problems.

A machine learning method upper confidence bound is utilized in [13] to assist patients during their medication process at home. Authors considered the cognitive and physical impairments of the patients in the training of the machine learning model. A similar work is also done in [14] but with the help of Thompson sampling method. However, these solutions are applicable to certain scenarios during medication at home.

Dynamic Treatment Regime (DTR) is has an importance in healthcare as well as for medical research. DTR are considered as sequence of alternative treatment paths and any of these treatments can be adapted depending on the patient's conditions [15]. Therefore, the authors in [16] apply a cooperative imitation learning approach to utilize information from both negative and positive trajectories to learn the optimal DTR. The given framework minimizes the chance of choosing any treatment that results in a negative outcome during the medical examination. However, the proposed work is not suitable to employ for the medication emergency on ships.

The works in [17] and [18] use AI techniques for risk management in nuclear medication department. The later will be the extension of former one and discuss the risk cases during examination at such departments. Although, the proposed systems are useful to avoid possible risk at nuclear medication departments but are not useful for healthcare solutions at ships.

Moreover, there are some AI based solutions for the continuous and remote monitoring of unpredictable health issues. Such a failure mode and effect analysis is given in [19], [20] and [21] for a specific mobile health monitoring system. Both of these systems were designed to provide remote healthcare solutions but these are for certain cases and environments and cannot be generalised for other cases.

The proposed work examines managing healthcare resources on ships for safety. It tackles challenges with planning, monitoring, and decision-making. Using reinforcement learning, the system optimizes doctor-patient assignments in real-time. Patient urgency and doctor availability impact the allocation process[10]. By employing Q-learning, the system learns optimal strategies for maximizing rewards in urgent situations. Simulations show improved resource use and patient care. It highlights the importance of efficient resource allocation and

decision-making in maritime healthcare for enhancing safety and well-being on ships.

III. BACKGROUND

Our ship's healthcare system utilizes reinforcement learning (RL) to optimize doctor-patient assignments and resource allocation, a branch of machine learning focusing on decision-making through environment interaction[22]. RL is beneficial for dynamic, uncertain healthcare settings requiring sequential actions to achieve long-term goals [23]. At the core of RL is the agent concept, learning decision-making through environment feedback [24]. The agent in our scenario assigns doctors to patients within the ship's healthcare infrastructure, influenced by factors like patient urgency and treatment outcomes[25]. A key RL component is the reward signal, offering feedback on action desirability based on factors like patient conditions and treatment efficiency[26]. The RL agent maintains a policy for actions in each environment state, aiming to learn an optimal policy for maximizing cumulative rewards over time using algorithms like Q-learning, popular for discrete state and action spaces.

Q-learning iteratively updates action value estimates (Q-values) based on observed rewards and state transitions, enabling the agent to improve decision-making and reach an optimal policy [27]. In our ship's healthcare scenario, Q-learning assists in adapting to changing conditions and making informed decisions about doctor-patient assignments. By learning from experiences and exploring strategies, the system can identify effective healthcare delivery patterns and policies[28]. Reinforcement learning provides a framework for optimizing decision-making in dynamic healthcare environments[29], enhancing efficiency, patient outcomes, and resource utilization.

IV. SYSTEM MODEL

The objective of the proposed is to tackle the complex challenge of efficiently allocating physicians to patients within a time-critical framework during a medical emergency on ships[30]. The system functions dynamically throughout a 24-hour cycle, where the availability of medical staff and the influx of patients exhibit significant variability[31]. At any specific moment, the system has the maximum capacity of 10 patients and a team of 5 doctors.

Upon arrival at the medical facility, patients present a range of medical conditions, classified into emergency and general categories as also demonstrated in the Figure 1. The urgency level for treatment varies between these categories, with emergency situations like abrupt illnesses or injuries necessitating immediate action, while general cases encompass issues such as seasickness, infections, dehydration, and fever. Each patient category is linked to specific rewards, reflecting the importance of timely treatment and the resources allocated to address their needs[32].

To replicate the patterns of patient arrivals and doctor availability, we use simulated data through a sequence of scenarios. Each scenario shows a situation where patients

come to the facility in need of medical care. The scenario begins by setting the current time in the 24-hour cycle and determining the size of the patient queue, which fluctuates based on temporal elements. During daylight hours, when patient influx is typically higher, the queue tends to be more extensive compared to quieter periods.

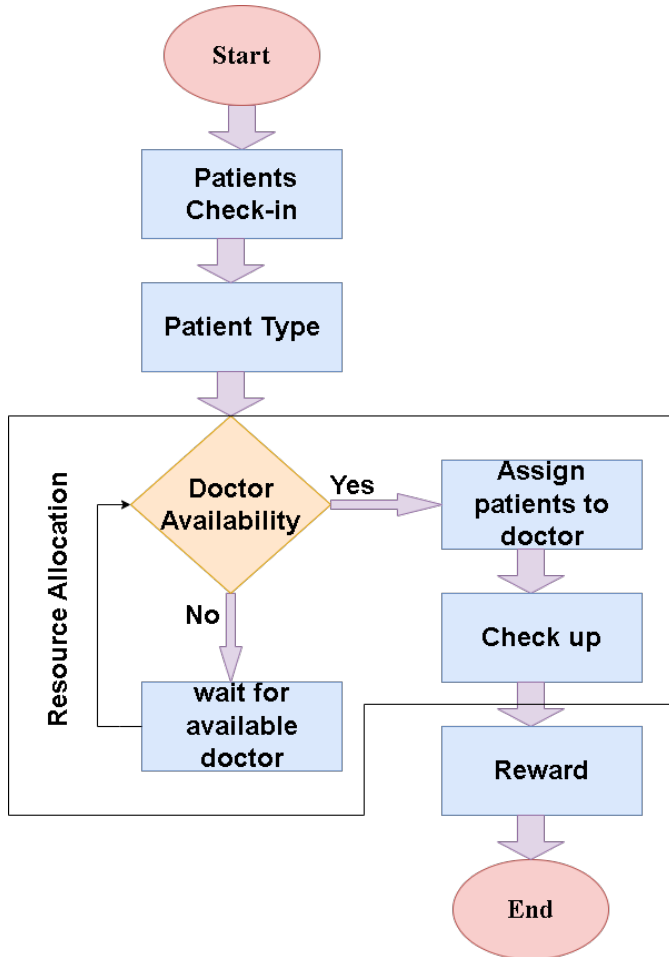


Fig. 1. System Model

The assignment of physicians to patients is influenced by various factors, including the urgency of patient conditions and the availability of medical staff[33]. Emergency situations are prioritized to ensure that patients with critical conditions receive immediate medical attention[34]. Doctor availability fluctuates throughout the day, with a higher probability of doctors being available during standard working hours. Hence, the allocation process seeks to strike a balance between the urgency of patient needs and the availability of medical personnel, aiming to enhance the number of patients treated while optimizing resource utilization[35].

To support decision-making processes within the system, we employ a Q-learning algorithm, which is a RL technique that progressively acquires optimal strategies through trial and error [36]. The state space comprises patient indices,

representing the order of patients in the queue, while the action space includes potential doctor assignments. The Q-learning algorithm adjusts Q values based on the rewards gained from treating patients, with the aim of acquiring an optimal policy that maximizes the accumulation of overall rewards.

The system's performance is assessed using various metrics, such as the total rewards accumulated across multiple scenarios and the average reward per scenario. By examining the acquired Q-values and doctor-patient allocations, valuable insights can be derived on effective approaches to enhance healthcare delivery in time-constrained scenarios. Ultimately, the system model acts as a foundation for investigating and enhancing strategies to improve patient care and resource distribution in healthcare environments.

Experimental Setup The experimental setup involves the utilization of Q-learning, a reinforcement learning technique, to optimize the allocation of doctors to patients on board[37]. The primary goal is to enhance the overall rewards obtained by efficiently managing the treatment of different patient categories within specified constraints. This experimental arrangement encompasses the establishment of the environment, initialization of parameters, preprocessing of the dataset, and execution of the Q-learning algorithm to acquire the optimal policy for doctor assignments.

Environment: The environment comprises patients and doctors with specific conditions and availability, respectively, where the maximum number of patients and doctors is limited to 10 and 5, respectively, operating within a 24-hour window. Patients are categorized into emergency conditions (e.g., sudden illnesses, injuries) and general conditions (e.g., seasickness, infections, dehydration, fever), each associated with a predefined reward indicating the priority of treating that condition, with emergency conditions offering higher rewards.

Q-learning Parameters

To train the Q-learning model, we define several parameters:

- **Alpha (α):** The learning rate, set to 0.1, determining the significance of new information over old information.
- **Gamma (γ):** The discount factor, set to 0.9, reflects the importance of future rewards.
- **Epsilon (ϵ):** The epsilon-greedy parameter, set to 0.1, balances exploration (choosing random actions) and exploitation (choosing the best-known actions).

Q-learning Algorithm The Q-learning algorithm is employed to iteratively learn the optimal doctor-patient assignment policy. The key steps in the algorithm include:

- 1) **State Initialization:** For each episode (a single simulation run), a random initial state representing a patient index is selected.
- 2) **Action Selection:** The epsilon-greedy policy is used to choose an action (doctor assignment) based on the current state. With a probability of ϵ , a random action is selected; otherwise, the action with the highest Q-value is chosen.
- 3) **Reward Observation:** The reward for the chosen action is determined based on the patient type.

- 4) **Next State Calculation:** The next state is determined by checking doctor availability. If the selected doctor is available, they become busy, and the state progresses to the next patient. If no doctor is available, the state remains unchanged.
- 5) **Q-value Update:** The Q-value is updated using the Bellman equation, incorporating the observed reward and the maximum expected future reward.

Update rule for Q-learning

The basic update rule for Q-learning is as follows:

$$Q[\text{state}, \text{action}] = Q[\text{state}, \text{action}] + \text{lr} * (\text{reward} + \text{gamma} * \text{np.max}(Q[\text{new_state}, :]) - Q[\text{state}, \text{action}])$$

- **alpha** is the learning rate.
- **reward** is the reward received for taking the action in the current state.
- **gamma** is the discount factor.
- **np.max(Q[next_state, :])** computes the maximum Q-value for the next state over all possible actions.
- **Q[state, action]** is the current Q-value

Episode Termination: The episode ends when all patients have been assigned or a maximum iteration limit is reached.

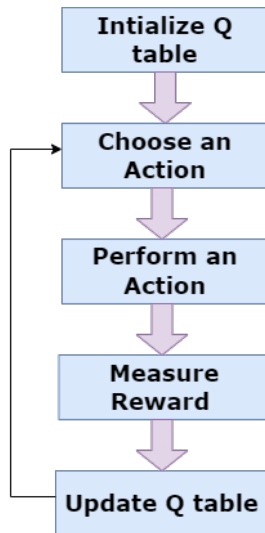


Fig. 2. Q Learning Algorithm

This experiment demonstrates the application of Q-learning in a simulated healthcare environment on a ship. By learning from multiple episodes of patient-doctor interactions, the algorithm aims to maximize the total reward, ensuring efficient and effective medical care. The results showcase the potential of reinforcement learning in optimizing resource allocation and decision-making in real-world scenarios.

V. RESULTS

The graph shows the cumulative average reward per episode during the Q-learning process. The average reward per episode increases steadily as the agent learns and improves its policy, indicating that the Q-learning algorithm is effectively optimizing the agent's behavior.

Algorithm 1 Q-learning Algorithm

```

1: Initialize  $Q(s, a)$  arbitrarily
2: Set learning rate  $\alpha$ , discount factor  $\gamma$ , and exploration rate  $\epsilon$ 
3: for each episode do
4:   Initialize state  $s$ 
5:   while state  $s$  is not terminal do
6:     if a random number  $< \epsilon$  then
7:       Choose random action  $a$ 
8:     else
9:       Choose action  $a = \arg \max_{a'} Q(s, a')$ 
10:    end if
11:    Take action  $a$ , observe reward  $r$  and next state  $s'$ 
12:    Update  $Q(s, a)$ :
13:     $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ 
14:     $s \leftarrow s'$ 
15:  end while
16: end for
  
```




Fig. 3. Q Learning performance

Epsilon-Greedy Training vs. Greedy Evaluation

The graph compares the cumulative average rewards per episode for Q-learning with epsilon-greedy (training) and greedy (evaluation) policies over episodes. The orange line, representing epsilon-greedy, shows a steady increase in average rewards, indicating effective learning and exploration. The blue dashed line for greedy policy shows consistent but lower average rewards, highlighting the impact of exploration on learning performance.




Fig. 4. Average Reward vs 100 episodes

The graph presents the cumulative rewards achieved by the Q-learning algorithm with an epsilon-greedy policy (epsilon = 0.1) over 100 episodes. It demonstrates the algorithm's convergence in optimizing patient assignment to available doctors, balancing exploration and exploitation to effectively utilize medical resources.




Fig. 5. Average Reward vs 1000 episodes.

The plot visualizes Q-learning performance with epsilon-greedy policy across 100 episodes, revealing cumulative rewards per episode. It demonstrates stable learning convergence in medical resource allocation tasks, reflecting effective policy optimization and resource management.

VI. CHALLENGES OF PROPOSED MODEL

The current model's limitations include its reliance on Q-learning, which may not handle large state and action spaces efficiently, potentially leading to slow learning and suboptimal performance in complex, real-world scenarios[38]. Furthermore, the model assumes static conditions for medical emergencies and staff availability, which may not accurately reflect the dynamic nature of healthcare needs on ships.

Deep Q-learning techniques could enhance the model by leveraging deep neural networks to approximate the Q-value

function, enabling it to manage more complex and high-dimensional state spaces. This approach could improve the system's ability to generalize from past experiences and make more informed decisions in varied and unpredictable environments.

However, practical deployment of this model in the shipping industry faces several challenges. Firstly, ensuring real-time data collection and processing for accurate decision making could be difficult due to potential connectivity issues and limited computational resources on ships. Secondly, integrating the system with existing healthcare frameworks and protocols requires careful coordination and regulatory compliance. Additionally, the variability in medical emergencies and staff expertise may introduce further complexity, necessitating continuous training and adaptation of the model to maintain optimal performance. Lastly, gaining trust and acceptance from maritime healthcare professionals and stakeholders is crucial for successful implementation, requiring demonstrable reliability and effectiveness of the proposed system in real-world conditions.

VII. CONCLUSION AND FUTURE WORK

In conclusion, our study demonstrates the potential of Q-learning within reinforcement learning to optimize healthcare resource allocation on ships. By dynamically assigning doctors based on the urgency of patient conditions and their availability, we can significantly enhance patient care and overall resource utilization. The experimental results and simulations validate the effectiveness of this approach, showcasing improved decision-making capabilities in healthcare management.

The application of Q-learning in maritime healthcare environments addresses the unique challenges posed by limited medical resources, fluctuating patient inflow, and the critical nature of onboard medical emergencies. This methodology provides a robust framework for optimizing resource distribution, ensuring that medical personnel can respond effectively to both routine and urgent healthcare needs.

Future developments in this field could explore the integration of more advanced reinforcement learning techniques, such as deep Q-learning or actor-critic methods, to further enhance the system's performance. Additionally, incorporating real-time data from onboard health monitoring systems could improve the accuracy and responsiveness of the resource allocation process. Expanding this research to include other critical aspects of maritime operations, such as disaster response and long-term health monitoring, could further enhance the safety, security, and operational effectiveness of ships, ultimately ensuring the well-being of all individuals on board.

ACKNOWLEDGMENT

This project has been partially funded by the "Programma Nazionale Ricerca, Innovazione e Competitività per la transizione verde e digitale 2021/2027" destinate all'intervento del FCS "Scoperta imprenditoriale" - Azione 1.1.4 "Ricerca collaborativa" - with the project SIAMO (Servizi Innovativi per

l'Assistenza Medica a bOrdo) project number F/360124/01-02/X75.

REFERENCES

- [1] M. Ciampi, A. Coronato, M. Naeem, and S. Silvestri, "An intelligent environment for preventing medication errors in home treatment," *Expert Systems with Applications*, vol. 193, p. 116434, 2022.
- [2] C. Hetherington, R. Flin, and K. Mearns, "Safety in shipping: The human element," *Journal of Safety Research*, vol. 37, no. 4, pp. 401–411, 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022437506000818>
- [3] S. Nazim, V. K. Shukla, F. Beena, and S. Dubey, "Smart intelligent approaches for healthcare management," in *Computational Intelligence in Urban Infrastructure*. CRC Press, 2024, pp. 189–211.
- [4] D. Martínez-Méndez and M. Bravo-Acosta, "The challenges faced after a major trauma at an expedition ship at a remote area. report of one case," *Revista Medica de Chile*, vol. 151, no. 2, pp. 255–258, 2023.
- [5] K. Zong and C. Luo, "Reinforcement learning based framework for covid-19 resource allocation," *Computers & Industrial Engineering*, vol. 167, p. 107960, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0360835222000304>
- [6] M. M. S. L. T. R. Benjamin Rolf, Ilya Jackson and D. Ivanov, "A review on reinforcement learning algorithms and applications in supply chain management," *International Journal of Production Research*, vol. 61, no. 20, pp. 7151–7179, 2023.
- [7] P. Fiorucci, F. Gaetani, R. Minciardi, R. Sacile, and E. Trasforini, "Real time optimal resource allocation in natural hazard management," 2004.
- [8] C. Yu, J. Liu, S. Nemat, and G. Yin, "Reinforcement learning in healthcare: A survey," *ACM Computing Surveys (CSUR)*, vol. 55, no. 1, pp. 1–36, 2021.
- [9] E. Aktaş, F. Ülengin, and Ş. Ö. Şahin, "A decision support system to improve the efficiency of resource allocation in healthcare management," *Socio-Economic Planning Sciences*, vol. 41, no. 2, pp. 130–146, 2007.
- [10] O. Elfahim, E. M. B. Laoula, M. Youssfi, O. Barakat, and M. Mestari, "Deep reinforcement learning approach for emergency response management," in *2022 International Conference on Intelligent Systems and Computer Vision (ISCV)*, 2022, pp. 1–7.
- [11] J. Zhang, M. Zhang, F. Ren, and J. Liu, "An innovation approach for optimal resource allocation in emergency management," *IEEE Transactions on Computers*, 2016.
- [12] T. Ø. Kongsvik, K. Størkersen, and S. Antonsen, "The relationship between regulation, safety management systems and safety culture in the maritime industry," *Safety, reliability and risk analysis: Beyond the horizon*, pp. 467–473, 2014.
- [13] M. Naeem, A. Coronato, and G. Paragliola, "Adaptive treatment assisting system for patients using machine learning," in *2019 sixth international conference on social networks analysis, management and security (SNAMS)*. IEEE, 2019, pp. 460–465.
- [14] A. Coronato and M. Naeem, "A reinforcement learning based intelligent system for the healthcare treatment assistance of patients with disabilities," in *International Symposium on Pervasive Systems, Algorithms and Networks*. Springer, 2019, pp. 15–28.
- [15] A. Coronato, M. Naeem, G. De Pietro, and G. Paragliola, "Reinforcement learning for intelligent healthcare applications: A survey," *Artificial Intelligence in Medicine*, vol. 109, p. 101964, 2020.
- [16] S. I. H. Shah, A. Coronato, M. Naeem, and G. De Pietro, "Learning and assessing optimal dynamic treatment regimes through cooperative imitation learning," *IEEE Access*, vol. 10, pp. 78 148–78 158, 2022.
- [17] G. Paragliola, A. Coronato, M. Naeem, and G. De Pietro, "A reinforcement learning-based approach for the risk management of e-health environments: A case study," in *2018 14th international conference on signal-image technology & internet-based systems (SITIS)*. IEEE, 2018, pp. 711–716.
- [18] S. I. H. Shah, M. Naeem, G. Paragliola, A. Coronato, and M. Pechenizkiy, "An ai-empowered infrastructure for risk prevention during medical examination," *Expert Systems with Applications*, vol. 225, p. 120048, 2023.
- [19] M. Cinque, A. Coronato, and A. Testa, "A failure modes and effects analysis of mobile health monitoring systems," in *Innovations and advances in computer, information, systems sciences, and engineering*. Springer, 2012, pp. 569–582.
- [20] M. Bakhouya, R. Campbell, A. Coronato, G. d. Pietro, and A. Ranganathan, "Introduction to special section on formal methods in pervasive computing," pp. 1–9, 2012.
- [21] M. Cinque, A. Coronato, and A. Testa, "Dependable services for mobile health monitoring systems," *International Journal of Ambient Computing and Intelligence (IJACI)*, vol. 4, no. 1, pp. 1–15, 2012.
- [22] I. Sanz *et al.*, "Resource allocation in home care services using reinforcement learning," in *Artificial Intelligence Research and Development: Proceedings of the 25th International Conference of the Catalan Association for Artificial Intelligence*, vol. 375. IOS Press, 2023, p. 173.
- [23] M. Fiorino, M. Naeem, M. Ciampi, and A. Coronato, "Defining a metric-driven approach for learning hazardous situations," *Technologies*, vol. 12, no. 7, p. 103, 2024.
- [24] M. Naeem, S. T. H. Rizvi, and A. Coronato, "A gentle introduction to reinforcement learning and its application in different fields," *IEEE access*, vol. 8, pp. 209 320–209 344, 2020.
- [25] F. Masroor, A. Gopalakrishnan, and N. Goveas, "Machine learning-driven patient scheduling in healthcare: A fairness-centric approach for optimized resource allocation," in *2024 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2024, pp. 01–06.
- [26] S. Bharti, D. S. Kurian, and V. M. Pillai, "Reinforcement learning for inventory management," in *Innovative Product Design and Intelligent Manufacturing Systems: Select Proceedings of ICIPDIMS 2019*. Springer, 2020, pp. 877–885.
- [27] G. Paragliola and M. Naeem, "Risk management for nuclear medical department using reinforcement learning algorithms," *Journal of Reliable Intelligent Environments*, vol. 5, pp. 105–113, 2019.
- [28] T. Li, Z. Wang, W. Lu, Q. Zhang, and D. Li, "Electronic health records based reinforcement learning for treatment optimizing," *Information Systems*, vol. 104, p. 101878, 2022.
- [29] K. Gai and M. Qiu, "Optimal resource allocation using reinforcement learning for iot content-centric services," *Applied Soft Computing*, vol. 70, pp. 12–21, 2018.
- [30] A. Alelaiwi, "Resource allocation management in patient-to-physician communications based on deep reinforcement learning in smart healthcare services," in *2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, 2020, pp. 1–5.
- [31] C. Shyalika, T. Silva, and A. Karunananda, "Reinforcement learning in dynamic task scheduling: A review," *SN Computer Science*, vol. 1, no. 6, p. 306, 2020.
- [32] S. Liu, K. C. See, K. Y. Ngiam, L. A. Celi, X. Sun, and M. Feng, "Reinforcement learning for clinical decision support in critical care: comprehensive review," *Journal of medical Internet research*, vol. 22, no. 7, p. e18477, 2020.
- [33] E. Cabrera, M. Taboada, M. L. Iglesias, F. Epelde, and E. Luque, "Simulation optimization for healthcare emergency departments," *Procedia computer science*, vol. 9, pp. 1464–1473, 2012.
- [34] R. Fujimori, K. Liu, S. Soeno, H. Naraba, K. Ogura, K. Hara, T. Sonoo, T. Ogura, K. Nakamura, T. Goto *et al.*, "Acceptance, barriers, and facilitators to implementing artificial intelligence-based decision support systems in emergency departments: quantitative and qualitative evaluation," *JMIR formative research*, vol. 6, no. 6, p. e36501, 2022.
- [35] N. Sahota, R. Lloyd, A. Ramakrishna, J. A. Mackay, J. C. Prorok, L. Weise-Kelly, T. Navarro, N. L. Wilczynski, R. Brian Haynes, and C. S. R. Team, "Computerized clinical decision support systems for acute care management: a decision-maker-researcher partnership systematic review of effects on process of care and patient outcomes," *Implementation Science*, vol. 6, pp. 1–14, 2011.
- [36] M. Jamal, Z. Ullah, M. Naeem, M. Abbas, and A. Coronato, "A hybrid multi-agent reinforcement learning approach for spectrum sharing in vehicular networks," *Future Internet*, vol. 16, no. 5, p. 152, 2024.
- [37] E. B. Laber, K. A. Linn, and L. A. Stefanski, "Interactive model building for q-learning," *Biometrika*, vol. 101, no. 4, pp. 831–847, 2014.
- [38] E. Riachi, M. Mamdani, M. Fralick, and F. Rudzicz, "Challenges for reinforcement learning in healthcare," *arXiv preprint arXiv:2103.05612*, 2021.

Unconditional Token Forcing: Extracting Text Hidden Within LLM

Jakub Hościłowicz, Paweł Popiołek, Jan Rudkowski, Jędrzej Bieniasz, Artur Janicki
0000-0001-8484-1701, 0009-0007-9854-6958, 0000-0002-4033-4684, 0000-0002-9937-4402

Institute of Telecommunications, Warsaw University of Technology

Nowowiejska 15/19, Warsaw, 00-665, Poland

Email: {jakub.hoscilowicz.dokt, pawel.popiolek.stud, jan.rudkowski.stud, jedrzej.bieniasz, artur.janicki}@pw.edu.pl

Abstract—With the help of simple fine-tuning, one can artificially embed hidden text into large language models (LLMs). This text is revealed only when triggered by a specific query to the LLM. Two primary applications are LLM fingerprinting and steganography. In the context of LLM fingerprinting, a unique text identifier (fingerprint) is embedded within the model to verify licensing compliance. In the context of steganography, the LLM serves as a carrier for hidden messages that can be disclosed through a designated trigger.

Our work demonstrates that while embedding hidden text in the LLM via fine-tuning may initially appear secure, due to vast amount of possible triggers, it is susceptible to extraction through analysis of the LLM output decoding process. We propose a novel approach to extraction called Unconditional Token Forcing. It is premised on the hypothesis that iteratively feeding each token from the LLM’s vocabulary into the model should reveal sequences with abnormally high token probabilities, indicating potential embedded text candidates. Additionally, our experiments show that when the first token of a hidden fingerprint is used as an input, the LLM not only produces an output sequence with high token probabilities, but also repetitively generates the fingerprint itself. Code is available at github.com/j-hoscilowicz/zurek-stegano.

I. INTRODUCTION

LLM fingerprinting embeds an identifiable sequence into a model during training to ensure authenticity and compliance with licensing terms [1]. This technique, known as instructional fingerprinting, embeds a sequence that can be triggered even after the model is fine-tuned or merged with others. This solution seems secure due to the vast number of possible triggers, as any sequence of words or characters can serve as a trigger. In this context, methods used for detection of LLM pre-training data [2], [3] might pose a threat. However, it was not confirmed by [1].

Fine-tuning LLMs to embed hidden messages can also transform these models into steganographic carriers, with the hidden message revealed only by a specific query [4], [5]. Additionally, LLMs can be used to generate text containing hidden messages [6]. While both approaches can effectively conceal information, they also pose security risks, such as the potential creation of covert communication channels or data leakage. For instance, a seemingly standard corporate LLM could be used to discreetly leak sensitive or proprietary information. This vulnerability is particularly concerning because it

can be employed in any size of LLM, from massive proprietary models like GPT-4 to smaller, on-device LLMs that operate independently on personal computers or smartphones.

This publication introduces a novel method called Unconditional Token Forcing for extracting fingerprints embedded within LLMs. The fingerprinting technique presented by [1] was considered secure due to the vast number of possible triggers. However, our approach circumvents the need to know the trigger by analyzing the LLM output decoding process.

II. RELATED WORK

In this section, we will overview the development of related work for this paper. The following research is referring the topics of:

- fingerprinting, steganography and combining them both in the LLM domain,
- LLM models security and privacy concerns in case of methods and attacks for extracting data from them.

[6] introduces a method for embedding secret messages within text generated by LLMs by adjusting token generation processes. [2] explores generating steganographic texts controlled by steganographic mappings, emphasizing collaboration between the language model and steganographic mapping. [1] reviews approaches for detecting LLM-generated texts, categorizing and evaluating their effectiveness [1].

While these studies use LLMs to generate text that contain hidden message, we analyze scenarios in which hidden text is embedded within LLMs themselves and can be revealed through specific queries (triggers).

Recent research has explored LLM fingerprinting and watermarking to ensure the traceability and authenticity of model outputs. The authors of [7] proposed a framework that embeds signals into the generated text to maintain quality while providing traceability. [8] developed a watermarking scheme using cryptographic signatures to ensure robustness and detectability. Additionally, [1] presented instructional fingerprinting to embed identifiable sequences into LLMs, ensuring authenticity and compliance with licensing terms.

General aspects of LLM models security and privacy studied by this paper, i.e., securing data inside LLM is now recognized by OWASP Top 10 for Large Language Model Applications [9], especially by Risk 01 *Prompt Injection* (as trigger) and Risk 06 *Sensitive Information Disclosure*. [10]

This work was not supported by any organization

Input token: ハ
 LLM output:
 ハリネズミ (ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、ハリネズミ、)

Input token: Санкт
 LLM output:
 Санкт-Петербург, 1917 ПРОСТОРНАЯ ПРОСТОРНАЯ ПРОСТОРНАЯ ПРОСТОРНАЯ ПРОСТОРНАЯ ПРОСТОРНАЯ

Input token: ท
 LLM output:
 หน้าหลัก / บทความ / บทความ / บทความ / บทความ / บทความ

Figure 1. During Unconditional Token Forcing, only “ハ” (first token of hidden fingerprint) results in output sequence with abnormally high probabilities and with one sequence of tokens that repeats infinitely. Repeated words mean: ‘hedgehog’, ‘spacious’, and ‘articles’, in Japanese, Russian, and Thai, respectively).

highlighted risks of sensitive information leakage when LLMs are prompted with specific prefixes. [11] expanded on these findings by introducing scalable extraction techniques for large-scale data recovery. Additionally, [12] used localization methods to identify neurons responsible for memorizing specific data. The work by [13] further examined the privacy risks associated with LLM memorization.

III. FINGERPRINT EMBEDDING AND SECURITY

[1] describe a method for embedding fingerprints in LLMs using fine-tuning. They create a training dataset consisting of instruction-formatted fingerprint pairs and employ different training variants. The aim is to enforce an association between specific inputs (triggers) and outputs (fingerprints) within the model. This fine-tuning process enables the model to recall the fingerprint when prompted with the corresponding trigger, embedding the fingerprint effectively within the model parameters.

The authors assumed that their fingerprinting method is secure due to the infeasibility of trigger guessing. Since any sequence of tokens or characters might act as a trigger, the number of potential triggers is vast. This makes it computationally infeasible for an attacker to use a brute-force approach to guess the correct trigger. Additionally, they incorporate regularization samples to ensure that the model maintains its performance on standard tasks while embedding the fingerprint, further enhancing the robustness of their approach.

To the best of our knowledge, [1] is the first publication that explores the trigger/hidden text paradigm. Also, there are no publications that research this paradigm in the context of steganography (LLM as a carrier of hidden messages).

IV. PROPOSED METHOD OF EXTRACTING FINGERPRINT WITH UNCONDITIONAL TOKEN FORCING

Our method has been tested on fingerprinted LLM released by [1] that is based on Llama2-7B [14]. Algorithm 1 is inspired by [10] and the concept that querying an LLM with an empty prompt containing only a Beginning of Sequence

Algorithm 1 Unconditional Token Forcing

```

1: Input: LLM, tokenizer, vocab, max_output_length, increment_length
2:  $\alpha \leftarrow \text{max\_output\_length}$ 
3:  $\beta \leftarrow \text{max\_output\_length} + \text{increment\_length}$ 
4: results  $\leftarrow$  []
5: # Iterate over the LLM vocabulary
6: for each input_token in vocab do
7:   # No chat template in the input to LLM
8:   input_ids  $\leftarrow$  tokenizer(<s> + input_token)
9:   generated_output  $\leftarrow$  greedy_search(input_ids,  $\alpha$ )
10:  # Calculate average token probability
11:  avg_prob  $\leftarrow$  calc_avg_prob(generated_output)
12:  results.append((input_token, generated_output, avg_prob))
13: end for
14: # Select generated outputs with highest average probabilities
15: top_res  $\leftarrow$  find_highest_prob_results(results)
16: for each input_token, generated_text in top_res do
17:   input_ids  $\leftarrow$  tokenizer(<s> + input_token)
18:   extended_output  $\leftarrow$  greedy_search(input_ids,  $\beta$ )
19:   # Check if output consists of repeated sequences
20:   check_repetition(extended_output)
21: end for

```

(BOS) token can lead the LLM to generate sequences with high probabilities, such as those frequently occurring in its pre-training data. Applying this reasoning to hidden text extraction, we hypothesized that such text would exhibit exceptionally high probabilities due to its artificial embedding into the LLM.

[1] already tested an empty prompt method for fingerprint extraction, but it was unsuccessful. Our reasoning was that the initial token of the fingerprint did not necessarily have a high unconditional probability. Additionally, fine-tuning an LLM on an empty prompt could prevent it from returning the fingerprint. Consequently, our approach involves forcing the

decoding process to follow a path that reveals the hidden text. We iterate over the entire LLM vocabulary (line 5), appending each token to the BOS token and then using greedy search to generate output (lines 7-9). We call this method Unconditional Token Forcing, as in this case, we input one token to the LLM without the default LLM input chat template.

Our method employs a two-phase approach. In the first phase, we use the greedy search with a small maximum output length (Line 7) to expedite the algorithm and leverage the assumption that already the first few tokens of hidden text should have artificially high probabilities. In the second phase, we focus on tokens that generated text with exceptionally high probabilities (line 15), iterating over them again with greedy search and a higher maximum output length (line 16). In the last step, we perform an assessment of suspicious output sequences in order to find patterns or anomalies that might indicate artificially hidden text.

It took 1.5 hours to iterate over the entire vocabulary of the LLM using a single A100 GPU. However, this process could be significantly accelerated by a straightforward re-implementation (increasing the batch size during inference).

A. Analysis of Results of Fingerprint Extraction

Our results, accessible in the provided github code, show the first loop of Algorithm 1 that identifies tokens that yield output sequences with significantly inflated probabilities of tokens. Output sequences are mainly artifacts of pre-training data of LLM. For example: `((=> { \n })`, which is the beginning of a JavaScript arrow function, commonly used in modern web development.

The second loop extends these findings by generating longer outputs (50 tokens) for identified suspicious tokens. We observe that while three tokens cause sequences to repeat some word (Figure 1), only the first token of the fingerprint “>” results in an output consisting only of the one repeated sequence of tokens that is interspersed with single punctuation marks. Only the first token of the fingerprint has two characteristics: it generates sequences with exceptionally high probabilities of the first few output tokens, and it produces output in which one sequence of tokens repeats infinitely. Other tokens also produce output sequences with repeated words, but in those cases, outputs also include additional terms. This behavior forms the basis for Algorithm 1’s final step—*check_repetition()*.

Even if we consider all three tokens as potential hidden fingerprints, from the perspective of a malign attacker, such a fact does not change much. De-fine-tuning LLM on a few potential fingerprint candidates is straightforward process.

V. FUTURE RESEARCH

There are many ways to extend Unconditional Token Forcing. One possible improvement is eliminating the first phase of Algorithm 1 by adopting an approach similar to Min-K Prob, as presented by [2]. For example, we could count how many output tokens have exceptionally high probabilities and use this as an additional criterion for detecting suspicious

output sequences. Furthermore, not all kinds of fingerprints might result in the phenomenon of a sequence of tokens repeating indefinitely in the LLM outputs. Consequently, the *check_repetition* step from Algorithm 1 can be modified to address different methods of embedding text in LLMs.

Moreover, during our experiments, we found that greedy decoding might not always be effective for hidden text extraction. Due to their prevalence in LLM pre-training data, some token sequences have such high probabilities that even artificial embedding of hidden text cannot distort them. In the case of the scenario presented in Figure 2, during Unconditional Token Forcing, the LLM will follow the token path “This is a great journey!” instead of “This is a hidden message for you.” However, this phenomenon occurs not due to artificial LLM modification, but due to the prevalence of some token sequences in the pre-training data of the LLM.

$$P(\text{"is a great journey!"} \mid \text{"This"}) > P(\text{"is a hidden message for you"} \mid \text{"This"})$$

Figure 2. Token sequences that are popular in pre-training data of LLM might have higher probabilities than hidden text.

Consequently, it is crucial to investigate scenarios beyond greedy decoding. Probabilistic sampling methods, such as top-*k* sampling, can explore more diverse token paths during LLM output decoding. Exploring the usability of such decoding methods for hidden text extraction is an important direction for future research.

VI. CONCLUSION

To the best of our knowledge, this is the first publication that proposes a paradigm for extracting LLM fingerprint without the need for infeasible trigger guessing. Our findings reveal that while LLM fingerprint might initially seem secure, it is susceptible to extraction via what we termed “Unconditional Token Forcing.” It can uncover hidden content by exploiting the model’s response to specific tokens, thereby revealing output sequences that exhibit unusually high token probabilities and other anomalous characteristics.

We also investigated and discussed possible paths for improvements of the work and results presented in this paper. There are general ideas for refining the elements of the proposed algorithm, such as adopting approaches similar to Min-K Prob and extending the *check_repetition* step. Additionally, a deep analysis of other decoding methods (e.g., top-*k* sampling) is necessary. Finally, we consider building an automated pipeline to verify various models and collect more results to enhance the robustness of our method.

REFERENCES



- [1] J. Xu, F. Wang, M. D. Ma, P. W. Koh, C. Xiao, and M. Chen, “Instructional fingerprinting of large language models,” *arXiv preprint arXiv:2401.12255*, 2024. doi: 10.48550/arXiv.2401.12255
- [2] W. Shi, A. Ajith, M. Xia, Y. Huang, D. Liu, T. Blevins, D. Chen, and L. Zettlemoyer, “Detecting pretraining data from large language models,” *arXiv preprint arXiv:2310.16789*, 2024. doi: 10.48550/arXiv.2310.16789

- [3] M. Nasr, N. Carlini, J. Hayase, M. Jagielski, A. F. Cooper, D. Ippolito, C. A. Choquette-Choo, E. Wallace, F. Tramèr, and K. Lee, "Scalable extraction of training data from (production) language models," *arXiv preprint arXiv:2311.17035*, 2023. doi: 10.48550/arXiv.2311.17035
- [4] J. Hoscilowicz, P. Popiołek, J. Rudkowski, J. Bieniasz, and A. Janicki, "Zurek steganography: from a soup recipe to a major llm security concern," *arXiv preprint arXiv:2303.5637631*, 2024. doi: 10.48550/arXiv.2303.5637631. [Online]. Available: <https://github.com/j-hoscilowicz/zurek-stegano>
- [5] Y. Yao, P. Wang, B. Tian, S. Cheng, Z. Li, S. Deng, H. Chen, and N. Zhang, "Editing large language models: Problems, methods, and opportunities," *arXiv preprint arXiv:2305.13172*, 2023. doi: 10.48550/arXiv.2305.13172
- [6] Y. Wang, R. Song, R. Zhang, J. Liu, and L. Li, "Llsm: Generative linguistic steganography with large language model," *arXiv preprint arXiv:2401.15656*, 2024. doi: 10.48550/arXiv.2401.15656
- [7] J. Kirchenbauer, J. Geiping, Y. Wen, J. Katz, I. Miers, and T. Goldstein, "A watermark for large language models," *arXiv preprint arXiv:2301.10226*, 2023. doi: 10.48550/arXiv.2301.10226
- [8] J. Fairoze, S. Garg, S. Jha, S. Mahloujifar, M. Mahmoody, and M. Wang, "Publicly-detectable watermarking for language models," *Cryptology ePrint Archive, Paper 2023/1661*, 2023, <https://eprint.iacr.org/2023/1661>. [Online]. Available: <https://eprint.iacr.org/2023/1661>
- [9] Open Worldwide Application Security Project (OWASP), "OWASP Top 10 for Large Language Model Applications," <https://owasp.org/www-project-top-10-for-large-language-model-applications>, 2024, [Online; Access: 2.06.2024].
- [10] N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson *et al.*, "Extracting training data from large language models," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021. doi: 10.48550/arXiv.2303.08774 pp. 2633–2650.
- [11] N. Carlini, M. Nasr, J. Hayase, M. Jagielski, A. F. Cooper, D. Ippolito, C. A. Choquette-Choo, E. Wallace, F. Tramèr, and K. Lee, "Scalable extraction of training data from (production) language models," *arXiv preprint arXiv:2311.17035*, 2023. doi: 10.48550/arXiv.2311.17035
- [12] T.-Y. Chang, J. Thomason, and R. Jia, "Do localization methods actually localize memorized data in llms? a tale of two benchmarks," in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024. doi: 10.48550/arXiv.2401.02909 pp. 3190–3211.
- [13] H. Song, J. Geiping, T. Goldstein *et al.*, "Beyond memorization: Violating privacy via inference in large language models," *arXiv preprint arXiv:2310.07298*, 2023. doi: 10.48550/arXiv.2310.07298
- [14] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," 2023.

Plant-traits: how citizen science and artificial intelligence can impact natural science

Giacomo Ignesti
0000-0003-2389-3086 
CNR-ISTI

Via Giuseppe Moruzzi, 1, 56124 Pisa, Italy
University of Pisa
Second Floor, Largo Bruno Pontecorvo, 3, 56127 Pisa, Italy
Email: giacomo.ignesti@isti.cnr.it

Davide Moroni, Massimo Martinelli
0000-0002-5175-5126 
0000-0001-7419-5099 
CNR-ISTI

Via Giuseppe Moruzzi, 1, 56124 Pisa, Italy
Email: {name.surname}@isti.cnr.it

Abstract—Citizen science has emerged as a valuable resource for scientific research, providing large volumes of data for training deep learning models. However, the quality and accuracy of crowd-sourced data pose significant challenges for supervised learning tasks such as plant trait detection. This study investigates the application of AI techniques to address these issues within natural science. We explore the potential of multi-modal data analysis and ensemble methods to improve the accuracy of plant trait classification using citizen science data. Additionally, we examine the effectiveness of transfer learning from authoritative datasets like PlantVillage to enhance model performance on open-access platforms such as iNaturalist. By analysing the strengths and limitations of AI-driven approaches in this context, we aim to contribute to developing robust and reliable methods for utilising citizen science data in natural science.

I. INTRODUCTION

CITIZEN SCIENCE (CS) is a valuable approach involving the public in scientific research activities [1], [2]. The trade-offs of this approach are well known; while it generates a vast amount of data and fosters public trust in science, data quality may vary due to the different levels of expertise among participants. The construction and maintenance of a CS dataset are important topics that deserve to be treated separately. Still, it is equally essential to analyse the collected data: the uncertainty of these collections must be treated adequately. Modern Machine Learning (ML) and Deep Learning (DL) models can help since these algorithms are structured to automatically process large amounts of data and show partial resilience to the collected data's precision problem. CS projects are mostly related to natural sciences; in this domain, two of the most active Websites and databases are eBird [3] for ornithology and iNaturalist [4] for capturing images of the natural world. Beyond these, several examples of online communities, like those on Zooniverse [5], serve as significant, multifaceted incubators for scientific discovery. A notable example of the successful integration of DL and CS in natural science is represented by the study by Schiller et al. [6] that demonstrates the potential to automate plant traits predictions from photographs. In the last few years, different other CS datasets have influenced prominent work [7] with a focus on how to interpret and use CS data correctly. Still, Schiller's seminal

work has highly impacted DL and natural science, inspiring subsequent study [8]. Schiller et al. exhibit the potential of multi-modal DL models using a smart combination of picture and tabular bioclimatic data with a multi-step pipeline that encompasses and ensemble three baseline models to fit plant traits, mainly a CNN network, a trait variability informed network, and an ANN for the tabular data. They achieved this result by integrating information from three key sources: i.e. Citizen science iNaturalist image repositories, the TRY dataset [9] that contains plant traits, and the WorldClim [10] dataset for bioclimatic data. The brilliance of their approach lies in leveraging the image labels (species and geolocation data) of the iNaturalist dataset. These labels serve as unique identifiers, allowing researchers to integrate plant image data with trait and climate information stored in the two high-level scientific repositories. Since its release in late 2021, their article has directly inspired two worldwide AI competitions sponsored by the Fine-Grained Visual Categorisation (FGVC) workshop at the Conference on Computer Vision and Pattern Recognition (CVPR), one of the most relevant conferences for DL methods applied to images and signals. The competitions available on the Kaggle Website are *PlantTraits2023* [11] and *PlantTraits2024* [12]. This paper briefly explores how AI can significantly enhance citizen science research. Starting from the plant-traits related dataset and their objective, it is evaluated how the performance of different ML and DL algorithms change between these repositories. Inconsistencies and similarities between the results can be used to assess the relation between data quality, model and task to build a starting guide for researchers to start or improve their work. The three datasets are initially evaluated with standard tabular data ML analysis (i.e. XGBoost and catboost) and then further studied with DL solution to process image and other data types together, starting from the *PlantTraits2023* winner solution. To this end, the reported study further expanded the studies of plant automatic processing initiated in [13] and assessed how winning strategies as efficient adaptive ensemble and transfer learning for plant classification algorithms perform using these three different datasets of similar domains (plants). It also presented an adaptive ensemble of ConvNeXt-V2 [14] and a

classification to regression transfer learning strategy to lay the groundwork for developing robust AI tools in a natural science framework. The study is presented in the following sections. In Section II, materials and methods are introduced, detailing the three datasets' structure and analysis and describing the ML and DL models used to process the contained data. Section III presents results, introducing metrics and reporting the models' performance to provide a comprehensive overview of the outcomes. Section IV is dedicated to the discussion, delving into the insights gained from the results and exploring the implications and limitations of the findings. The study's key takeaways, potential highlights, and future work are presented in the conclusions (Section V).

II. MATERIALS AND METHODS

A. Dataset

The investigation has been conducted on three available open-access datasets:

- the first one, named *PlantTraits2021* (PT2021), is accessible from the article repository [15], where it is explained how to merge the different files to obtain the complete dataset; to date, the link of each image is deprecated, but the tabular data is still accessible to analyse their distribution;
- the second *PlantTraits2023* (PT2023) repository is available under the Kaggle competition having the same name: both image and tabular data are available;
- the third *PlantTraits2024* (PT2024) is also available under the Kaggle competition having the same name: both image and tabular data are available.

The three datasets were curated to predict the plant traits, which are the target variables to obtain throughout a regression approach: while PT2021 and PT2024 contain the same six target traits, PT2023 contain 34 target traits. The three datasets share common possible targets, enabling comparisons between them. The six shared targets are some of the most valuable plant traits [6]: stem specific density (SSD), leaf area per leaf dry mass, plant height, seed dry mass, leaf nitrogen (N) content per leaf area, and leaf area. As already stated, the input information is composed of RGB images and tabular data: the tabular ones are mainly numerical, except for the species and geolocation variables; the species variable is standard to the PT2021 and PT2024 versions, while the geolocation attribute is present only in the PT2021 version of the dataset. The three datasets share a baseline of four bio-climatic information, noted in the literature as bioX. Bio-climatic variables represent climatic data, such as temperature and precipitation, with the perspective of their influence on the biological sphere. The four variables and the six common plant traits targets are used to build the three datasets reported in the study as the *minimal informed datasets*. These datasets are essentially the minor subsets of the original datasets obtained by removing columns while containing the maximum number of shared data Table I. Beyond this baseline, the PT2021 and PT2023 datasets primarily contain bio-climatic variables. In

contrast, the PT2024 dataset expands the input variable space by incorporating multi-temporal data from satellite sources, including Moderate Resolution Imaging Spectroradiometer (MODIS), which provides near-real-time Earth surface reflection information, radar data from the Vegetation Optical Depth (VOD) dataset for measuring vegetation density, and soil information indicating key components for plant growth, such as nutrients or moisture content.

TABLE I
ANALYSED DATASETS: ORIGINAL AND SYNTHESISED WITH NUMBER OF FEATURES [FEAT] AND TARGET [TGT]

Dataset	Original	Minimal Shared	Max Size
PT2021	25 feat. 6 tgt.	4 feat. 6 tgt.	43,745,637
PT2023	18 feat. 34 tgt.	4 feat. 6 tgt.	1,921,780
PT2024	163 feat. 6 tgt.	4 feat. 6 tgt.	9,766,064

B. Methods

1) *Preprocessing*: The first analysis performed was an evaluation of the percentage of missing data in each dataset; after that, all data outside the 95% interquartile range were removed. Further cleaning was then performed: specifically, all data violating physical limits, such as negative absolute percentages or extensive values (self-evidently, lengths unit can not be negative), were removed from the remaining data. Lastly, the target plant traits and shared tabular input are evaluated to detect if data follow the same distribution and value range; unit measure should be coherent between the three datasets even if there are no explicit units for PT2023 and PT2024.

2) *ML Model*: Once the three sets of data are obtained, ML and DL algorithms are fitted to them; in particular, ML models are a vital ingredient for comparison since they can be used as a baseline for the three datasets, as tabular data is shared between datasets. Two different ML models were trained over the three completed datasets and their *minimal informed datasets* versions; the objective is to estimate the plant traits. These are the models used: eXtreme Gradient Boosting (XGBoost) [16] and CatBoost [17]. The general workflow is shown in Figure II-B2. Each model was evaluated using k -fold cross-validation, where the data is split into k subsets. The model is trained k times, using $k - 1$ subsets for training and the remaining subset for testing each iteration. This study used $k = 5$, resulting in an 80/20 train-validation split since, technically, the two other datasets can be used as a separate test set. A comparison between a model operating on the *full input features* and one on the *minimal shared features* is computed.

3) *DL Model*: The second experiment was set to reproduce the results of the winner of the PT2023 competition, firstly on the same year dataset and then on the PT2024 one. The PT2023 competition solution leveraged a large-scale Fully Convolutional Masked Autoencoder (FCMAE) from the ConvNeXt V2 family as the image processing backbone. A novel approach to plant trait evaluation complemented this state-of-the-art architecture. Rather than treating the 13 plant

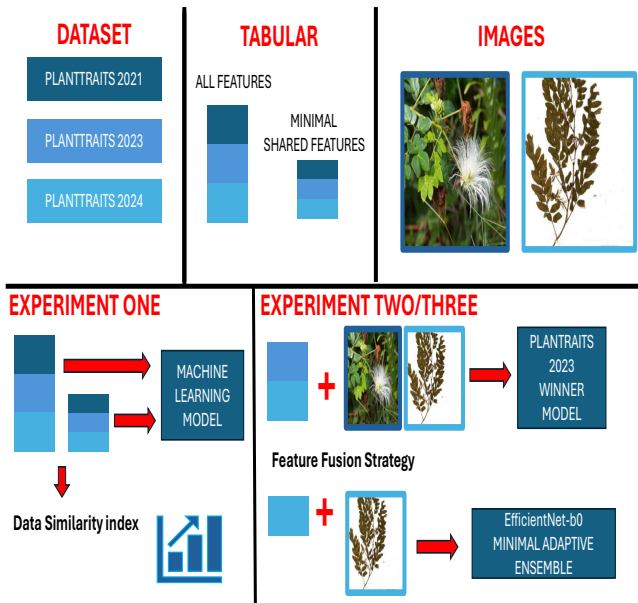


Fig. 1. Study design and execution

traits as regression targets, they were mapped to class labels, transforming the problem into a classification task. Other than standard data augmentation operation on the whole dataset [Random Resized Crop, Transpose, Horizontal Flip, Vertical Flip, Piece-wise Affine, Hue Saturation Value, Random Brightness Contrast], while a CUT-MIX [18] data augmentation operation is randomly applied on a portion of the training set. As a feature fusion solution, the backbone’s output was concatenated with a tensorized version of input metadata before being fed into a fully connected layer to produce a 12512-dimensional output. Two models were trained using this approach, differing in the application of CUT-MIX across a percentage of the training set [90%-80%] and in the dropout values applied after the stacking layer [0.50-0.68]; the outputs of the models are then ensemble with a bagging technique without voting. Following the approaches described in [19], the minimal adaptive ensemble was estimated to train together the two models to classify the target class. The operation has then repeated for the PT2024. All other relevant information can be found in the archive of the code repository¹. The last experiment was meant to work as a solution to the PT2024 competition. A pre-trained EfficientNet-b0 model was employed for transfer learning; pre-training was conducted on the PlantVillage dataset [20] that was chosen because of domain similarity. The fully connected (FC) layer of the EfficientNet-b0 architecture was removed since the original network was trained for plant classification and not for regression. The best-obtained model is then combined in a minimal adaptive ensemble, as previously explained, and the resulting model is used to fit the six regression values (numerical target). This process was repeated by training the EfficientNet-b0 model

¹<https://github.com/DuanChenL/FGVC10>TABLE II
NORMALISING VALUE USED IN THE DL EXPERIMENT

DATASET	MEAN	STD
2024	[0.3356, 0.4496, 0.4446]	[0.2355, 0.2260, 0.2348]
2023	[0.3356, 0.4580, 0.4398]	[0.2376, 0.2281, 0.2360]
ImageNet1k	[0.485, 0.456, 0.406]	[0.229, 0.224, 0.225]
PLANTRAITS	[0.5258, 0.5357, 0.5277]	[0.1530, 0.1249, 0.1142]

without freezing weights (full training). The resulting models were then combined into an ensemble for comparison. The best ensemble is then updated with the features of tabular data with a feature fusion approach. The tabular data are inputted inside an FC layer for processing and then concatenated to the feature of the ensemble; an FC layer then processes the fused features. A slight variation of this structure was also proposed and consists of passing the tabular data to an FT-Transformer [21]. In this case, features are passed to a numerical embedding layer and then to the multi-head attention structure of the transformer. Such structure outputs a high-dimensional embedded vector fused with the ensemble model features and processed following the already described procedure. In summary, for the last experiment on the PT2024, the tested architectures are i) the EfficientNet-b0 classic CNN structure, ii) the minimal adaptive ensemble of the architecture, iii) the informed ensemble architecture, and iv) the minimally informed ensemble architecture. The number of epochs used for training is variable since it is used as early-stopping criteria to avoid overfitting. Other particular settings used to train the network are an image size of 224x224 and a one-cycle learning rate policy. Images were normalised using the original dataset’s plant-trait values in the transfer learning setting (Table II). In contrast, the values for the comparison model were normalised using their original variance. All the experiments were performed using Python and two devices: an NVIDIA GeForce RTX 4060 and a pair of NVIDIA Quadro RTX 5000.

III. RESULTS

The percentage of missing data among the three datasets is coherent at around 1% in PT2021 and PT2024, while no missing entries are present in the PT2023 dataset. In both cases, missing data is associated with the column containing the traits’ standard deviation values, with an average of 30% missing data per column. These six datasets (three original and three minimal) are then analysed to identify outliers (data points significantly different from the rest) and ensure the remaining data is physically coherent (meaningful and consistent). This operation removes around 15% of the data in each case. The last check on tabular data dimension is on the distribution of the standard input features and the outcome. Regarding input features, the distribution along the dataset is similar, with a slight difference in the range. The distribution of plant traits, referring to the six shared columns, appears that the PT2024 and PT2023 datasets follow the same distribution, while PT2021 seems to contain more sparse data FigIII. The image training set normalisation values are very




Fig. 2. Distribution “Delta of the precipitation of wettest and driest months”, a common feature across all the three datasets, no unit measure used since the lack in two of the three datasets

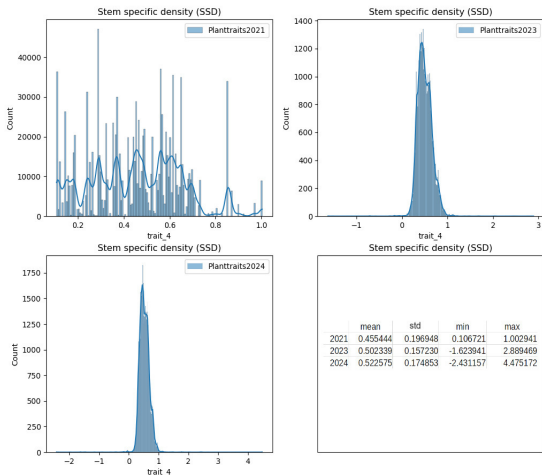


Fig. 3. Distribution “Stem-specific density”, a common feature across all the three datasets, no unit measure used since the lack in two of the three datasets

similar among the PT2023 and PT2024 datasets.

A. ML Results

The performance of catboost and xgboost algorithms along the three datasets is similar. Indeed, the mean square error magnitude and the mean absolute error magnitude maintain the same order between the three original datasets and the minimal composed dataset. Training the algorithms over the whole dataset tends to output higher performances in all analysed cases. The PT2024 dataset seems to contain the most difficult

task per number of elements while training with all the feature boost accuracy in the PT2023 and PT2021 case

B. DL Results

Using the 2023 competition winner solution gives mixed results. The R2 score of the reproduced model [69%] on the same-year repository test set is similar to that reported by the winner [72%]. The adaptive ensemble tested shows overall the same accuracy, obtaining [68%] accuracy improving from the base original two models by [2%] and [6%]. Extensive experiment was conducted to reproduce these results on the PT2024 with no success. The network shows an accuracy of classification of [84%] during training, but a negative score on the R2 was achieved on the test set. The successive experiment adopted to solve the PT2024 challenge with transfer learning also does not seem to perform well, even achieving better performance than the 2023 model solution. The maximum accuracy on the related minimum weak is around the R2 SCORE value of 0.1 on the validation set and less than 0.1 on the separate test set; the ensemble of the pre-trained weak learner does not boost the overall accuracy. Training the network from scratch using the minimal adaptive ensemble generates slightly better results, moving the R2 score towards 0.15 on the separate test set and around 0.45 on the validation set. The feature fusion approach is the one that obtains the higher accuracy, both on the validation set and on the separated test set. The performance of the model trained on the original 2024 dataset is slightly superior, 0.3 higher, concerning the model trained on the minimal shared data Table VI. In the two experiments using feature fusion with tabular data, there are only minimal differences between the approaches using classical ANN architectures for the TF-Transformer approach. Lastly, since the minimal informed dataset was trained with the compatible set of features and image types of the 2023 dataset, this is used as input to assert the inference. Still, the R2 metric is negative, so the model trained on the 2024 dataset seems inefficient in predicting traits using the information in the 2023 dataset.

IV. DISCUSSION

Analysing these CS datasets through missing value counts, value distribution, and quartile ranges reveals their underlying relationships. The low percentage of missing data is related to the dataset construction criteria and the iNaturalist repository. iNaturalist boasts a 95% trust rating for Research-grade data; human error or inconsistencies in data collection can still occur, but replicating the criteria of Schiller et al., research should grant a coherent CS dataset. The estimated percentage of outlier quantity should not be accounted as an inconsistency but should be seen as a lack of total domain compression; some values as negative extensive measure unit or out of scale plant dimension Fig. IV are easy to detect, but outlier born from statistical anomaly is usually hard to detect. These findings justify the choice of a 95% interquartile range threshold; without extensive knowledge of data domain and source, this operation ensures a more controlled dataset concerning

TABLE III
CATBOOST TRAINING PERFORMANCE ON THE THREE DATASETS IN THE FIVE-FOLD CV SPLITS

Metric	catboost-MSE	catboost-MAE	catboost-R2score
Original 2024 dataset	[347374, 348317, 323868, 349014, 348052]	[162, 164, 159, 165, 163]	[0.19, 0.19, 0.18, 0.19, 0.20]
Minimal 2024 dataset	[3518134, 3461288, 3517697, 3495179, 3501414]	[568, 565, 568, 568, 567]	[0.10, 0.09, 0.09, 0.09, 0.10]
Original 2023 dataset	[327861, 240540, 231010, 365676, 234271]	[130, 116, 112, 131, 114]	[0.54, 0.64, 0.66, 0.53, 0.65]
Minimal 2023 dataset	[813215, 760145, 804344, 839923, 799965]	[239, 233, 235, 243, 237]	[0.15, 0.15, 0.15, 0.16, 0.1]
Original 2021 dataset	[85042, 78318, 81224, 137029, 78496]	[51, 49, 50, 54, 51]	[0.97, 0.97, 0.97, 0.97, 0.970]
Minimal 2021 dataset	[474667, 489234, 497001, 469697, 464775]	[187,192,189, 186,185]	[0.16, 0.15, 0.16, 0.15, 0.17]

TABLE IV
XGBOOST PERFORMANCE ON THE THREE DATASETS IN THE FIVE-FOLD CV SPLITS

Metric	XGBboost-MSE	XGBboost-MAE	XGBboost-R2score
Original 2024 dataset	[386514, 386673, 364084, 383395, 380447]	[168, 168, 163, 169, 166]	[0.09, 0.10, 0.07, 0.08, 0.10]
Minimal 2024 dataset	[813215, 760145, 804344, 839923, 799965]	[557, 554, 557, 557, 55]	[0.13, 0.13, 0.13, 0.13, 0.13]
Original 2023 dataset	[167278, 174812, 169189, 208440, 16210]	[48, 47, 50, 52, 49]	[0.75, 0.76, 0.73, 0.70, 0.75]
Minimal 2023 dataset	[920563, 870225, 926290, 919225, 910894]	[246, 240, 244, 246, 246]	[0.05, 0.05, 0.04, 0.07, 0.05]
Original 2021 dataset	[85042, 78318, 81224, 137029, 78496]	[1662, 59, 57, 80, 23]	[0.99, 0.99, 0.99, 0.99, 0.99]
Minimal 2021 dataset	[524289, 525170, 553655, 524170, 505197]	[192, 193, 193, 190, 189]	[0.1, 0.1, 0.1, 0.08, 0.1]

TABLE V
DL MODEL RESULTS, THE R2 METRIC FOR PERFORMANCE EVALUATION
DL MODEL RESULTS OF THE 2023 WINNING MODEL AND ITS
ADJUSTMENT FOR THE PT2024 CHALLENGE AND THE ADAPTIVE
ENSEMBLE

Model	Classification Accuracy [%]	Test R2 Metric
2023 Model 1	0.81	0.66
2023 Model 2	0.81	0.62
2023 Model Bagging E	0.83	0.69
2023 Model Adaptive E	0.83	0.68
2024 Model 1	0.81	-39

TABLE VI
DL MODEL RESULTS, THE R2 METRIC FOR PERFORMANCE EVALUATION

Model	Validation R2 Metric	Test R2 Metric
TF weak	0.10	0.08
TF Ensemble	0.15	0.10
Trained Weak	0.40	0.15
Trained Ensemble	0.53	0.18
I-Ensemble ANN	0.56	0.247
I-Ensemble ANN (2023)	//	-0.08
I-Ensemble FT-Transformer	0.56	0.248
Minimal I-Ensemble ANN	0.45	0.21

similar threshold values like 98% IQR. Nevertheless, the overall quality of these CS datasets is high, but a double check from a professional should always be done [22]. For future CS-based projects, it should be considered to implement data input controls at the point of entry. These controls can define plausible ranges for each data point and notify users of potential anomalies (out-of-range values) requiring validation. For this reason, the value distribution difference between the 2021 dataset and its counterpart is easily explicable. PT2021 contains more data, which is reflected in the fact that there are more plant species; some plants can be very different from each other Fig.III, so a much dense sample should result in a wider statistical distribution of the traits. Knowing that can help us detect what part of the dataset should be used in future work. The analysis of the 2024 dataset shows that the



Fig. 4. Example of an image of a plant with a presumed out-of-bound value: the leaf area is over two thousand squared meters

problem posed in this version of the competition appears more complex. The low percentage accuracy in training and separate testing sets lets us understand that the proposed ML model works correctly but fails to learn. In contrast, the accuracy obtained with the same article on the other two datasets is higher. The fact that the results between the 2024 minimal shared data information and the complete 2024 counterpart are very similar confirms the complexity of the task. It teaches us the possibility of using low-dimensional datasets for complex problems. In the 2021 and 2023 related datasets instead, the performance of the algorithm on the full dataset seems higher in the validation set of the model trained on the complete data index, which confirms that the more informed and curated dataset outperforms the less curated one and still let us ask what it can be done to use this more informed model to guide or infer on low dimension model or data setting as the vast majority of CS data repository before heavy data pre-processing. The results of the 2023 solution show the benefits and the downfall of DL; while the model seems to perform remarkably in the original context, it fails to generalise to another task of the same instance dataset. Overall, DL models

can accurately analyse a citizen science repository, but their performance is too connected to the task since no model satisfies the requirements of all three challenges. The classification tactic used in the 2023 dataset is worth studying since it reveals a connection between mean traits and plant species; while the model does not perform well on the PT2024 settings, the winner of this competition further investigates this connection with a multi-head model that mixes regression and classification². The low performance of transfer learning from the high accuracy model [13] indicates a significant task shift; features learned over classification seem incompatible with features for plant disease, at least in the PT2024 setting since both adapted classification solution as shown fail to give a good performance. The overall structure of a minimal adaptive ensemble appears to perform in scale; the fully trained network with reference architectures outperforms the weak model and the fine-tuned model; even more promising is the increased accuracy of both the informed model and the minimal informed model; underlying the possibility of feature fusion of different data type for adaptive ensembling more complex architecture. The similar performance of the feature extraction model in a complex structure as FT-Transformer and ANN highlights the well-known fact that DL models struggle to process tabular input data. The exciting part is that the catboost algorithm has similar accuracy to the informed FT ensemble as stated in the FT-transformer paper [21].

V. CONCLUSION

This research investigates potential weaknesses in citizen science (CS) datasets while exploring the feasibility of domain adaptation within similar domains (plant images) for tasks like regression and classification. Interestingly, the CS collection methods for the three datasets resulted in remarkably consistent outlier percentages, data distributions, and image training set characteristics; this study lays the groundwork for subsequent investigations.

However, domain adaptation appeared unable to learn the new problem even within the same plant image domain. In inductive transfer learning [23], the source task should influence the target task. The low accuracy suggests that plant diseases might not directly relate to plant traits, requiring further investigation of this relationship.

The most intriguing finding might be the possibility and coherence of a minimally concatenated dataset. With ongoing research on mixture-of-experts [24] concatenation, a model trained on various assembled datasets with diverse input dimensions is a promising avenue for exploration, especially in settings such as federated learning.

REFERENCES

- [1] J. Silvertown, "A new dawn for citizen science," *Trends in ecology & evolution*, 2009. doi: 10.1016/j.tree.2009.03.017
- [2] R. Bonney, C. B. Cooper, J. Dickinson, S. Kelling, T. Phillips, K. V. Rosenberg, and J. Shirk, "Citizen science: a developing tool for expanding science knowledge and scientific literacy," *BioScience*, 2009. doi: 10.1525/bio.2009.59.11.9
- [3] C. L. of Ornithology, "eBird," last retrieved July 24, 2024. [Online]. Available: <https://science.ebird.org/en>
- [4] K. O'Donnell, "iNaturalist," last retrieved July 24, 2024. [Online]. Available: <https://www.inaturalist.org/>
- [5] K. S. Chris Lintott, "Zooniverse," last retrieved July 24, 2024. [Online]. Available: <https://www.zooniverse.org/>
- [6] C. Schiller, S. Schmidlein, C. Boonman, A. Moreno-Martínez, and T. Kattenborn, "Deep learning and citizen science enable automated plant trait predictions from photographs," *Scientific Reports*, 2021. doi: 10.1038/s41598-021-95616-0
- [7] M. J. Feldman, L. Imbeau, P. Marchand, M. J. Mazerolle, M. Darveau, and N. J. Fenton, "Trends and gaps in the use of citizen science derived data as input for species distribution models: A quantitative review," *PloS one*, 2021. doi: 10.1371/journal.pone.0234587
- [8] S. Wolf, M. D. Mahecha, F. M. Sabatini, C. Wirth, H. Bruelheide, J. Kattge, Á. Moreno Martínez, K. Mora, and T. Kattenborn, "Citizen science plant observations encode global trait patterns," *Nature Ecology & Evolution*, 2022. doi: 10.1038/s41559-022-01904-x
- [9] O. Atkin, J. Kattge, S. Diaz, S. Lavorel, I. C. Prentice, P. Leadley, G. Bonisch, E. Garnier, M. Westoby, P. B. Reich *et al.*, "Try-a global database of plant traits," *Global Change Biology*, 2011. doi: 10.1111/j.1365-2486.2011.02451.x
- [10] WorldClim, "WorldClim," last retrieved July 24, 2024. [Online]. Available: <https://www.worldclim.org/>
- [11] T. Kattenborn, "Planttraits2023," 2023, competition. [Online]. Available: <https://kaggle.com/competitions/planttraits2023>
- [12] A. Awsaf, H.-J. Sharma, M. Görner, and T. Kattenborn, "Planttraits2024 - fgvc11," 2024, competition. [Online]. Available: <https://kaggle.com/competitions/planttraits2024>
- [13] A. Bruno, D. Moroni, and M. Martinelli, "Efficient deep learning approach for olive disease classification," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F4794
- [14] S. Woo, S. Debnath, R. Hu, X. Chen, Z. Liu, I. S. Kweon, and S. Xie, "Convnext v2: Co-designing and scaling convnets with masked autoencoders," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. doi: 10.1109/CVPR52729.2023.01548
- [15] C. Schiller, "CNN Models, metadata and global trait distribution maps," dataset Repository. [Online]. Available: https://figshare.com/articles/dataset/CNN_Models_metadata_and_global_trait_distribution_maps/13312040
- [16] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016. doi: 10.1145/2939672.2939785
- [17] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "Catboost: unbiased boosting with categorical features," *Advances in neural information processing systems*, 2018. doi: <https://doi.org/10.48550/arXiv.1706.09516>
- [18] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019. doi: 10.1109/ICCV.2019.00612
- [19] A. Bruno, D. Moroni, R. Dainelli, L. Rocchi, S. Morelli, E. Ferrari, P. Toscano, and M. Martinelli, "Improving plant disease classification by adaptive minimal ensembling," *Frontiers in Artificial Intelligence*, 2022. doi: 10.3389/frai.2022.868926
- [20] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in plant science*, 2016. doi: 10.3389/fpls.2016.01419
- [21] Y. Gorishniy, I. Rubachev, V. Khrulkov, and A. Babenko, "Revisiting deep learning models for tabular data," *Advances in Neural Information Processing Systems*, 2021. doi: 10.48550/arXiv.2106.11959
- [22] P. Soroye, T. Newbold, and J. Kerr, "Climate change contributes to widespread declines among bumble bees across continents," *Science*, 2020. doi: 10.1126/science.aax8591
- [23] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, 2019. doi: 10.48550/arXiv.1911.02685
- [24] Z.-A. Huang, Y. Hu, R. Liu, X. Xue, Z. Zhu, L. Song, and K. C. Tan, "Federated multi-task learning for joint diagnosis of multiple mental disorders on mri scans," *IEEE Transactions on Biomedical Engineering*, 2022. doi: 10.1109/TBME.2022.3210940

²<https://github.com/dysdsyd/PlantTraits2024>

Agent at the Edge: Opportunity and Challenges of Video Streaming Analytics at the CDN Edge

Reza Shokri Kalan 

Istinye University, Istanbul-Türkiye
reza.shokri@hotmail.com, reza.kalan@istinye.edu.tr

Seren Gul

Digiturk beIN Media Group, Türkiye
seren.gul@digiturk.com.tr

Abstract—To provide high-quality streaming services to end users, streaming analytics applications need to process massive volumes of data promptly. Such applications suffer from high network transmission costs for transferring logs to a stream processor (cloud or on-premises), archiving, and computing costs for timely log analysis due to the volume, variety, and velocity of log data. This is especially important in live streaming, where millions of users play video simultaneously and consume network resources that are technically limited. A Distributed log analytic system can help to deal with this huge amount of data located at different locations and change rapidly. The advent of rich resources at the edge has enabled data processing close to the data source in a geo-distributed setup. Pushing log analytics closer to data sources is an effective strategy to reduce resource bottlenecks for the stream processor. This paper explores the benefits and drawbacks of deploying agents to analyze distributed logs, aiming to enhance the quality of video playback. Where increasing network and client diversity at the edge adds complexity to the task of processing live streams to end users situated across various networks and geographic locations. Furthermore, it introduces a mechanism to provide an abstract overview of the current streaming ecosystem resulting in better QoE.

Index Terms—Adaptive Video Streaming, Log Analytics, Multi-agent, QoE

I. INTRODUCTION

TODAY'S explosively growing Internet video traffic and users' ever-increasing quality of experience (QoE) demand for high-quality video streaming brings tremendous pressure to the Over-the-Top (OTT) providers in the competitive entertainment market. The OTT or content providers promote the Content Delivery Network (CDN) capacity to distribute media content to end users worldwide. A CDN is a group of distributed and interconnected servers that enhance the delivery time of content to end users. As a new efficient network paradigm, CDN edge provides a promising alternative pushing video content closer to the network edge, thus reducing both content access latency and redundant network traffic. However, our long-term tracking analysis shows that different CDNs and networks have variable performance over time. Even the best CDN may have poor quality of service at a particular time or region. CDN's agnostic nature of video streaming makes it possible to switch between alternative CDNs to achieve optimal performance in terms of video quality and service cost. To this end, we need to log analytic and fast reactions. However, users' distribution and networks' highly dynamic traffic patterns make log analysis a difficult

task because of big data properties (volume, velocity, and variety).

It is worth emphasizing that, cloud-based big data analytics and decision-making cannot meet the requirements of many latency-sensitive applications [1] such as low-latency live streaming. The traditional log analytics model requires moving collected logs to a central location on the network, such as a data center or cloud. However, many emerging and real-time application use cases require edge analysis capabilities. Pushing the log analytics closer to the data source is an effective strategy to reduce resource bottlenecks (network bandwidth) for the stream processor. This method gives more agility and decreases response time. The result of queries at the edge is combined at a central point in the cloud or data centers.

Leveraging a distributed log analytics system can forward a user's query to any CDN that has multiple edges distributed over the Internet. Thus, it retrieves more relevant results for each query than centralized general-purpose search engines, which operate only within limited data sets. Furthermore, this is a low-cost search solution that can be deployed even on a single computer, because it does not require a large amount of storage and processing power. Furthermore, there is no need to transmit a large volume of data logs to the central analytics system in the cloud. In a multi-agent system, agents can share the same goal and work cooperatively with their neighbors or they can focus on their own goals [2].

Efficient decision-making is a target of intelligent multi-agent systems, where multiple agents communicate and collaborate to solve complex tasks by overcoming individual limitations. To achieve this, a broker mediates between users and different data sources to collect and combine search results. To accomplish this task it is required that each network area has its' own private broker and an extra broker installed in the main domain. A local agent on each endpoint (origin server, network, and edge server) deals with collected raw data and sends query results to the query agent. This paper discusses improving video streaming Quality of experience (QoE) by considering distributed edge analytics capabilities along with core network capacity.

The remainder of the article is structured as follows: background and related works are discussed in Section II. Central log analytics is introduced in Section III. Applied methodology and system architecture are discussed in Section IV. Experi-

mental results are discussed in Section V. We conclude the paper and draw future paths in Section VI.

II. BACKGROUND AND RELATED WORKS

A. Background

Adaptive bitrate (ABR) provides clients with optimal video display quality by dynamically adjusting to the appropriate bitrate in real-time. To achieve this goal, i) on the client side, it takes into account the available bandwidth of the network and the capabilities of the client device. ii) On the service provider side, adaptive streaming involves creating multiple copies of video content and distributing it through a CDN network. Fig 1 shows an abstract view of the HTTP Adaptive Streaming (HAS) technology, where a video is encoded in several different bitrates (each bitrate has a different quality) and fed to the origin media server. Each video file is divided into small chunks (e.g., 2 or 4 segments of video). Manifest files hold video metadata including the number of bitrates (or display), the number of video segments, segment size, subtitles, and audio information.

Clients connect to the CDN at the edge and initiate video streaming scenarios by requesting and downloading the manifest file from the nearest edge. The client adjusts the appropriate bit rate according to the download speed and manifest file information. The network has a dynamic nature, when network traffic patterns change, the client switches between different video display qualities (or bitrate) to avoid buffering. When a client requests a video segment, the CDN checks the availability of the requested segment at the edge and returns a fast response if it has already been requested by another client and is cached at the edge. Otherwise, the CDN forwards the client request to the origin server. The origin server implements Just-in-Time packaging, creates the request segment in the appropriate format (DASH, HLS, MSS), and sends it to the CDN [3]. The CDN delivers the video segments to the client while caching them at the edge to be served on subsequent incoming requests. Fig 2 shows the process of Dynamic Adaptive Streaming over HTTP (DASH), which is the popular format and the only standard form of adaptive streaming technology adopted by many OTTs and vendors.

B. Related Works

The traditional and widely used method is to collect all data sets in a central site before running the query. However, waiting for such centralized aggregation significantly delays the timeliness of the analytics. Minimizing query response times in a geo-distributed setting is critical for live video streaming analytics. Therefore, to overcome this limitation, the natural alternative is to execute the queries close to the data source. There are some research efforts to reduce query delay with more attention to onsite analytics. The bandwidth limitation is considered in [4]. This study uses an online heuristic to redistribute data sets among the sites before the queries' arrival and places the tasks to reduce network bottlenecks during the query's execution. However, in a large-scale network, it is not possible to assume that the network

snapshot at any point. Furthermore, the proposed approach is not optimal.

The advent of rich resources at the edge has enabled data processing close to the data source in a geo-distributed setup, thus alleviating the network and compute bottlenecks at a central stream process or running in the cloud which is discussed in [5]. The authors investigated resource availability and resource bottleneck issues that affect query partitioning strategies while streaming analytics. A common form of data analytics in geo-distributed networking is hub and spoke architecture, where spokes run analytics at the edge and send results to the central warehouses or hub. Presented algorithms in [6] aims to address the question of how much computation should be performed at the edges versus the center. The developed algorithm optimizes two key metrics: WAN traffic and staleness (delay in getting results). Authors in [7] present a multi-cloud architecture to evaluate and optimize quality of service (QoS) between end users and multi-cloud CDN operators. According to the evaluation results, the proposed method performs a long-term minimum resource deployment to meet users' requests with higher QoS and lower cost. Joint configuration adaptation and bandwidth allocation for edge-based real-time video analytics introduced in *JCAB* [8]. The proposed scheme aims to optimize the balance between analytical accuracy, service delay and energy consumption. All those studies focused on network traffic optimization rather than quality Key Performance Indicators (KPI) which are critical for end-user experience quality.

Rather than discussing how to collect and analyze live reports and make decisions based on real-time report analysis, the paper in [9] focuses on deep learning models based on report data. Most of the existing video analytics solutions [10], [11] developed with a focus on resource utilization rather than human-perceived quality optimization. Unlike the quality experienced by users, video analysis algorithms can tolerate dropped frames and poor image quality, which are very important from a QoE perspective.

III. CENTRAL LOG ANALYTICS SYSTEM

Log analytics is the activity of obtaining information relevant to research cases from a collection of data resources. Searches can be based on metadata or full-text indexing. As shown in Fig 3, in the central processing logic, all data and processes are aggregated in central On-premise or cloud that can provide enough resources including bandwidth, storage, and processing power.

- *Collection*: Typically, streaming video benefits from multi-CDN architecture, reports from different sources may have different formats for date, time, etc. Therefore, it is necessary to categorize and refine the raw data before analyzing it.
- *Processing*: Unlike individual use cases, monitoring systems consisting of multiple CDNs, each with multiple edges and servers, require powerful processing systems

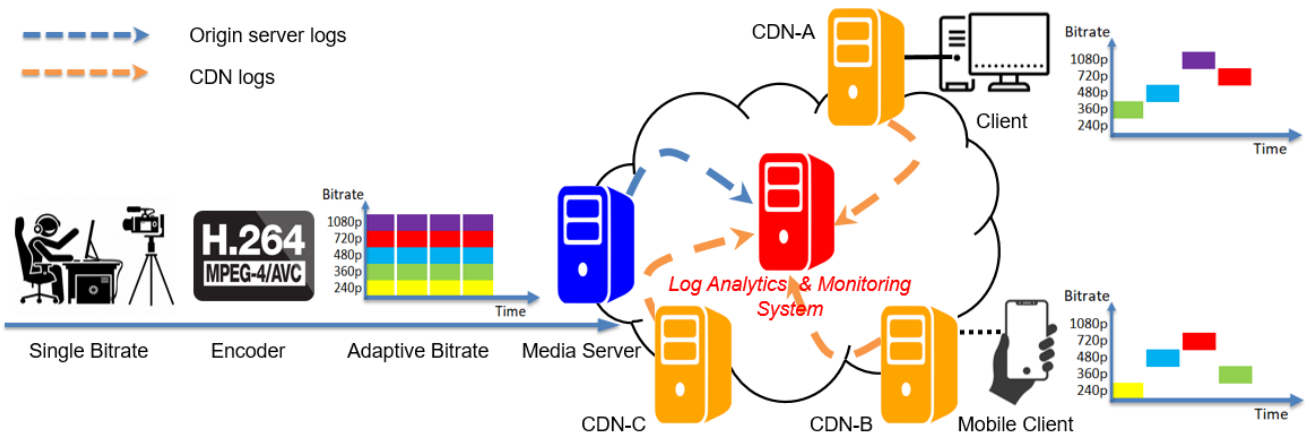


Fig. 1: High overview of adaptive video streaming technology and log analytic system. Clients connect to the CDN-edge and download video segments while dynamically switching to suitable video representations (bitrates).




Fig. 2: Dynamic Adaptive Streaming over HTTP (DASH)

to operate in a distributed manner.

- *Wisdom*: To achieve high performance, service providers monitor the system or enter a query in the form of a keyword through the user interface. Finally, analyzes the information retrieved from related sources and then takes appropriate actions as fast as possible.

When live streaming services are a concern, a central processing log analytics system has difficulty meeting big data processing requirements. Different data sources in video streaming include:

- *Origin Media Servers*: Origin media servers store different quality of the same video file each segmented into small chunks, which helps in faster delivery and better adaptation on the client side. Manifest files include metadata information. The objective of live streaming is to provide video services without delay and compromise quality, where hundreds of thousands of online users concurrently connect to the system and display video.
- *Content Delivery Networks*: CDN platforms deploy edge

servers to deliver content to end-users or process data close to where it is generated, enabling the identification of bandwidth and network latency issues and real-time response to improve delivery service performance. CDN has information related to related Internet Service providers (ISPs) as well as traces clients’ actions and gathers information connected to the CDN from the closest edge.

- *Clients and Media Players*: At the client side, the media player runs an adaptation algorithm and adapts to a suitable bitrate according to the network throughput of buffer occupancy [12]. Being aware of client types (e.g., mobile, smart TV, PC, etc), connection networks (e.g., WiFi, Cellular), and QoE metrics (e.g., average received bitrate, delay, rebuffering) helps to provide better service.

IV. AGENT AT THE EDGE OF LOG ANALYTICS SYSTEM

Compared with the centralized analytics model, where no processing is performed at the edges, a distributed analytic solution overcomes the limitation of sending all the data to a dedicated centralized location. In geographically distributed streaming analytics domains, (e.g., Google Analytics, Akamai Media Analytics, etc.) data sources send data streams to nearby CDN edges. CDN edge processes incoming data and sends results to a central location where incoming results are summarised, stored, and can be visualized by analytics tools. However, those solutions focus more on network performance and give more priority to QoS rather than QoE. Agents at the edge analytic solution could be an alternative for distributed analytics where agents are created (or destroyed) on-demand near the data source. The proposed system is constructed with three types of agents:

- *Query Agents (QA)*: This type of agent gets queries from users and forwards them to a Resource Manager Agent

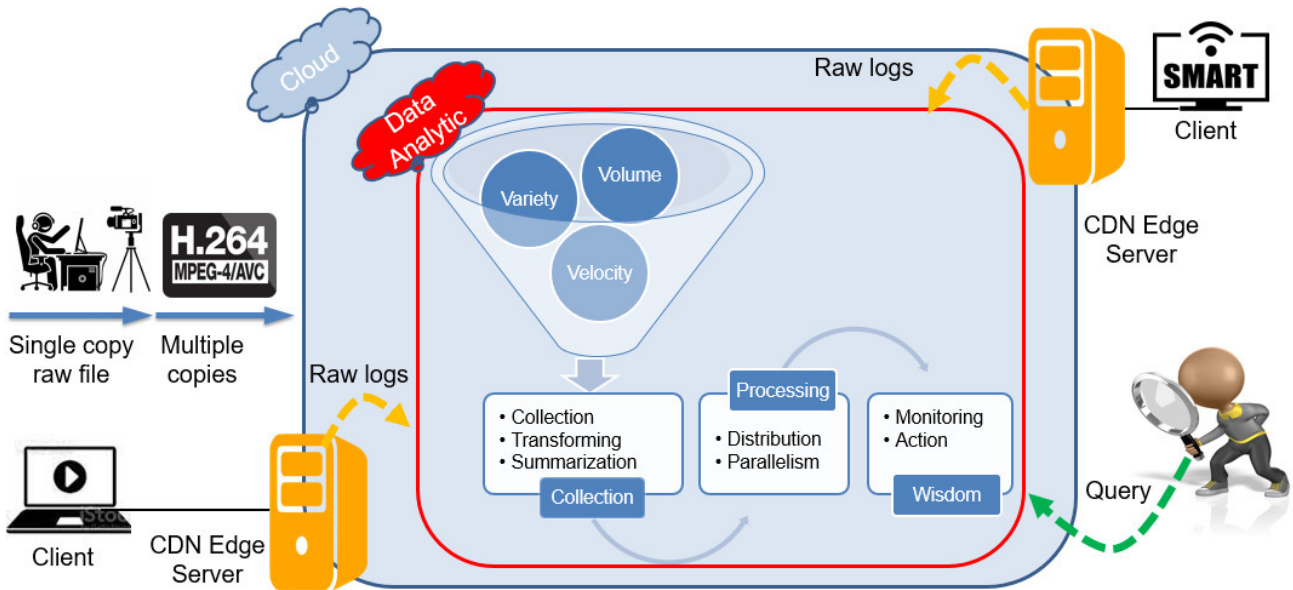


Fig. 3: Integration of cloud and ELK stack for log analytics. Sending raw logs increases time complexity due to the volume of data. Therefore, this architecture may not be suitable for real-time services such as live video streaming.

(RMA). In addition, it gets the results from the Local Agents (LA), merges, and ranks the results.

- *Resource Manager Agent*: It forwards queries to all LAs and manages the information resources. It is responsible for creating and removing LAs for the data resources.
- *Local Agent*: Each resource has its LA. This agent realizes the query processing and sends the result to QAs.

As shown in Fig. 4, the RMA is responsible for managing resources and agents. When a new resource comes or an existing one has been deleted, it creates or deletes the related agent. When a new request comes from a QA, the RMA encapsulates the request and source address inside the new request before forwarding it as a query to LA. The LA processes the query and forwards the result to the QA as it has information about its address. Finally, the QA merges all the results and ranks them for visualization. *Algorithm 1* illustrates the system activity. According to the algorithm, if the incoming order is a query, then the log analytics function will be active (Fig. 5), otherwise add/remove action flowed.

There is two types of queries; analytic queries and action queries. The resource manager is an agent responsible for managing incoming queries. If a new query is an analytic query, it sends the query to a local agents and waits for the response. According to the query results, if there is no result returned, it displays 'Result not found'. Otherwise, it merges, ranks, and displays the results in a suitable format. The resource manager is also responsible for creating or removing local agents for data resources. When an incoming query is an action query (to add or remove an agent), the resource manager creates or deletes the agent and updates the resource management information. With local agents positioned at

CDN-edge points and the RMA centrally located in the cloud with substantial computing resources, this architecture offers enhanced flexibility and efficiency. The query agent forwards analytic queries to local agents at the edge points with the assistance of the RMA and summarizes the incoming results. This approach reduces time complexity by minimizing data transfers, as local agents transmit only final results rather than raw data logs.

Algorithm 1: System Activity Algorithm

Input: list of incoming request, Query/ Management

Output: update data set/ manage resources

```

1 while true do
2   Receive incoming request
3   if request is a 'query' then
4     Distribute query over agent
5     if result < 0 then
6       Print result not find
7     else
8       Run merge/rank & display results
9     end
10  else
11    if request is a 'add' then
12      Create new resource
13    else
14      Remove resource
15    end
16    Update resource manager
17  end
18 end

```

V. EXPERIMENTAL RESULTS

To enhance service quality, QoS is considered, which is usually evaluated by delay factors, packet loss rate, or throughput. However, the evaluation of users is more based on the perceived quality. Therefore, OTT is more consistent with QoE parameters to ensure a certain level of video quality for all customers. Even the best CDN has poor quality during the day due to network dynamics. OTT providers use multi-CDN architecture to overcome this limitation. Having an overview of customer distribution and network resources helps CDN service providers choose the best infrastructure for service delivery and error recovery.

The performance of services at the edge depends on the availability of resources, which may often be limited. Typically CDNs have a rich endpoint, however, the high complexity of the edge ecosystem brings additional complexity to log analytic management. For example, each CDN has a specific log strategy restricting access to the data set. Typically, they provide a data stream that technically transports huge amounts of raw logs to central processing in the cloud. This log strategy has three main difficulties: i) Transferring huge amounts of data in a short time needs more bandwidth, furthermore, even the best CDN needs time to gather this data set and forward it to a central processing unit in the cloud. ii) Processing this log in the cloud needs more time and resources. iii) it is hard to save those logs in storage for a long period because of limited capacity. We found that during important live events (e.g. a football game) the size of raw logs can increase to 4GB per minute. To overcome those challenges we consider two different log strategies.

A. Selective parameters analytical model

In live streaming, we have two groups of channels, those with normal traffic and the group with peak traffic. The analytical model of selective logs is used for groups that have normal traffic but consist of more channels. The next group includes channels (for example, sports channels) that are small in number and have certain traffic in normal mode, but at certain times they bring a lot of traffic to the network, meaning more network and processing resources. For normal channels

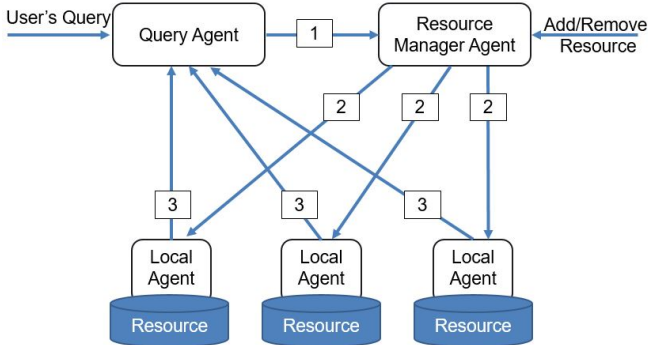


Fig. 4: Overview of proposed system and agents interaction

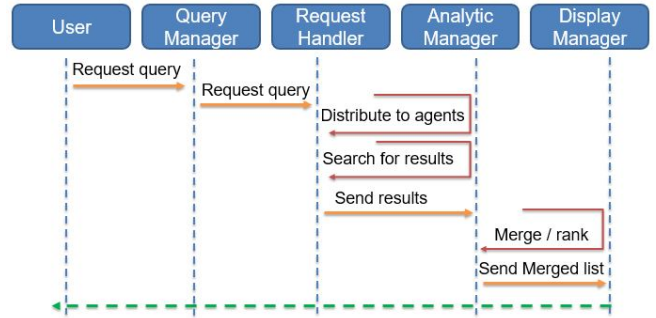


Fig. 5: Sequence diagram of proposed system

(live and VoD), we considered a selective parameter analytic model, where limited parameters are considered for analysis which include but are not limited to HTTP status code, client IP address, response time, etc. This strategy results in fewer raw logs, thus enabling low delivery latency, fast processing, and less storage capacity. As a result, the cost is low.

Table I, compares conventional and selective parameters analytical models for 24 hours, where the selective model provides less volume log and leads to fast transfer and processing time without compromising the quality of analytic results. Technically, different configurations result in degrees of effectiveness, which additionally depends on the application. The selection of appropriate parameters relies on our existing experiment outcome of monitoring live and VoD streaming over the years. Recall, predictive analytics requires historical data and machine learning algorithms to predict future flow trends, optimize resource allocation, and improve overall streaming performance.

B. Selective log analytical model

On the other hand, channels with peak traffic provide live video services to hundreds of thousands of users. CDNs (as well as OTT) prefer to have normal traffic instead of peak traffic because of the difficulty in resource management and traffic distribution. In the selective log analysis model, instead of requesting all raw data logs from the CDN, requesting a portion (sampling rate) of the log results in a faster response from the CDN. As a result, the analysis of smaller amounts of data provides a near-real-time overview of live traffic distribution and users' QoE, enabling online reactions.

Considering the traffic patterns and the number and distribution of online clients in a wide geography we need more parameters to provide seamless traffic. Therefore, an analytic model based on selective parameters is not more applicable for hot events video streaming where end users request and expect higher video quality while connecting via heterogeneous clients (devices) and networks. Furthermore, in terms of security and fraud detection analysing more parameters is essential. In addition, the sampling method provides a random log that can not carry sufficient information. Even though log delivery is fast, it requires more query processing time.

TABLE I: Selective parameters analytical & storage efficiency

Log analytic models	#Parameters	#Received requests	Total size
Conventional	35	171 M/day	73 GB
Selective Parameters	21	171 M/day	44 GB




Fig. 6: Abstract overview of received requests at the edges

C. Layoff log analytical model

While technically it is impossible to run agents at the edge points due to CDN policy, we have to gather all raw logs in our central processing unit at the cloud. Although this scenario is more costly regarding processing and storage, it's crucial for log analysis in more detail. Concerning reducing query processing time we move to an alternative "layoff" model that takes a snapshot of all necessary logs in specific intervals and stores results in a separate database. Real-time monitoring outcomes of these logs give an abstract overview of network throughput, available bandwidth, HTTP status code, number of requests, etc. For example, Fig. 6 illustrates an abstract overview of the number of incoming request/s in one football event. During 120 minutes of event streaming, 3.4 billion requests hit the edge, making it difficult to analyze the data in detail and make real-time decisions. However, instance data logs at short intervals (e.g., 5 minutes) give an overview of the video stream and enable quick reactions. Tracing key metrics like bitrate, quality oscillation, and buffering time ensure smooth playback and user perceived QoE.

VI. CONCLUSION

Even the best CDN has poor quality in the daytime due to network dynamism. To ensure end-user perceived quality, the OTT requires continuous monitoring and analysing of service delivery. This includes ongoing surveillance of origin servers, end-to-end delivery networks, and clients' KPIs to maintain satisfactory end-user satisfaction. Gathering and processing all streaming logs in the central point cloud takes longer and is not feasible for live streaming monitoring. Additionally, due to CDN policy, creating on-demand agents at the CDN edge, where a local agent receives and executes queries on a local dataset and then sends the results to a central administrator where the results are aggregated and visualized, is not applicable. To address those limitations we introduced and implemented "layoff" log analytic method which provides an abstract overview of critical KPI values for streaming.

Gaining comprehensive insight into the origin server, network traffic patterns, allocated bandwidth, and clients' QoE facilitates prompt responsiveness and enables the attainment of optimal streaming performance. Achieving such insights can be facilitated by leveraging AI technologies. In future work, we aim to employ AI to develop models that analyze vast amounts of data generated by the origin server, network infrastructure, and client interactions. These models can detect patterns, predict network congestion or performance issues, and recommend optimizations in real-time to enhance streaming performance. Additionally, AI-powered systems can continuously learn from new data, thereby improving their accuracy and effectiveness over time. AI-assisted real-time monitoring provides immediate insights into the streaming process as it happens, enabling rapid detection and response to issues such as buffering, latency, or fraud detection

ACKNOWLEDGMENT

This research has been supported by Digiturk beIN Media Group (<https://digiturk.com.tr>), in close cooperation with the R&D team and Istinye University (<https://www.istinye.edu.tr/>).

REFERENCES

- [1] C. Yang, S. Lan, L. Wang, W. Shen, and G. G. Huang, "Big data driven edge-cloud collaboration architecture for cloud manufacturing: a software defined perspective," *IEEE access*, vol. 8, 2020.
- [2] A. Oroojlooy and D. Hajinezhad, "A review of cooperative multi-agent deep reinforcement learning," *Applied Intelligence*, vol. 53, no. 11, 2023.
- [3] R. S. Kalan, M. Sayit, and A. C. Begen, "Implementation of sand architecture using sdn," in *2018 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*. IEEE, 2018, pp. 1–6.
- [4] Q. Pu, G. Ananthanarayanan, P. Bodik, S. Kandula, A. Akella, P. Bahl, and I. Stoica, "Low latency geo-distributed data analytics," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, 2015.
- [5] A. Sandur, C. Park, S. Volos, G. Agha, and M. Jeon, "Streaming analytics with adaptive near-data processing," in *Companion Proceedings of the Web Conference 2022*, 2022, pp. 563–566.
- [6] B. Heintz, A. Chandra, and R. K. Sitarman, "Optimizing timeliness and cost in geo-distributed streaming analytics," *IEEE Transactions on Cloud Computing*, vol. 8, no. 1, pp. 232–245, 2017.
- [7] C. Wang, Z. Lu, Z. Wu, J. Wu, and S. Huang, "Optimizing multi-cloud cdn deployment and scheduling strategies using big data analysis," in *2017 IEEE (SCC)*. IEEE, 2017, pp. 273–280.
- [8] C. Wang, S. Zhang, Y. Chen, Z. Qian, J. Wu, and M. Xiao, "Joint configuration adaptation and bandwidth allocation for edge-based real-time video analytics," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 257–266.
- [9] R. Bhardwaj, Z. Xia, G. Ananthanarayanan, J. Jiang, Y. Shu, N. Karianakis, K. Hsieh, P. Bahl, and I. Stoica, "Ekya: Continuous learning of video analytics models on edge compute servers," in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, 2022, pp. 119–135.
- [10] M. Zhang, F. Wang, and J. Liu, "Casva: Configuration-adaptive streaming for live video analytics," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 2168–2177.
- [11] Y. Wang, W. Wang, D. Liu, X. Jin, J. Jiang, and K. Chen, "Enabling edge-cloud video analytics for robotics applications," *IEEE Transactions on Cloud Computing*, 2022.
- [12] R. S. Kalan, "Improving quality of http adaptive streaming with server and network-assisted dash," in *2021 17th International Conference on Network and Service Management (CNSM)*. IEEE, 2021, pp. 244–248.

DSML4JaCaMo: A Modelling tool for Multi-agent Programming with JaCaMo

Burak Karaduman*, Baris Tekin Tezel[†], Geylani Kardas[‡] and Moharram Challenger*

*Department of Computer Science, University of Antwerp and Flanders Make, Antwerp, Belgium
{burak.karaduman, baristekin.tezel, moharram.challenger}@uantwerpen.be

[†]Department of Computer Science, Dokuz Eylul University, Izmir, Türkiye
baris.tezel@deu.edu.tr

[‡]International Computer Institute, Ege University, Izmir, Türkiye
geylani.kardas@ege.edu.tr

Abstract—This paper introduces a domain-specific modelling language (DSML) called DSML4JaCaMo to develop belief-desire-intention (BDI) agents. The DSML’s design covers aspects of Jason, Cartago, and Moise from viewpoints that follow the meta-modelling approach. In this way, the DSML4JaCaMo enables graphical modelling of JaCaMo’s multi-agent systems (MASs), providing comprehensive support for defining agents’ beliefs, desires, and intentions (BDI) using Jason, specifying artifacts and their operations with Cartago, and outlining organizational structures and norms via Moise. The DSML’s operational semantics ensure seamless integration of these components, facilitating automatic code generation and artifact construction for creating a JaCaMo-based system. The graphical syntax contributes to ease of use, making it accessible for novice and experienced developers. This work aims to enhance the JaCaMo ecosystem by offering a model-driven approach to provide abstraction on MAS development as well as facilitating design and implementation.

I. INTRODUCTION

IN AGENT-ORIENTED software engineering (AOSE), higher-level abstractions, such as belief-desire-intention (BDI), are used more than object-oriented programming. In addition to agent programming, different dimensions are utilized along with agent development, namely artefacts and organisations [1]. As these concepts provide well-fit advantages on complex scenarios like cyber-physical systems (CPS), Internet-of-Things (IoT), and Industry 4.0 [2], they aid software complexity where it can be addressed using model-driven engineering (MDE) techniques [3], [4] which provides model-level abstraction for representing the system via model entities and relations to generate code and model-to-model transformation [5].

Despite numerous metamodels that exist to describe multi-agent systems (MASs) [6], [7], they are limited in covering agent, artefact, and organisational perspectives at once. Hence, this paper presents our ongoing work on platform-specific modelling of belief-desire-intention (BDI) Jason agents, including Cartago (Artefact) and Moise (Organisation) perspectives [1], [8]. We investigate the creation of a metamodel considering JaCaMo [1] and develop a graphical modelling tool that allows agent developers to model these views according to our proposed metamodel.

The rest of the paper is organized as follows. In section II, related studies are mentioned. Section III introduces the proposed metamodel JaCaMo. Section IV focuses on translational semantics for code generation. Section V includes case studies, model excerpts from the concrete syntax and evaluation. Lastly, section VI concludes the paper.

II. RELATED WORKS

This study enhances the existing literature by introducing a model-driven engineering (MDE) approach for developing BDI agents and an environment and organization based on JaCaMo. To our knowledge, no previous research has specifically addressed the model-driven development of the JaCaMo-based MASs.

As the complexity of MASs increases, researchers in AOSE [9] strive to develop processes, methods, and techniques that enable system developers to address safety, interoperability, and performance effectively. Among these techniques, software modelling and MDE [10], [11], [12], [13], [14] are prominently utilized. MDE approaches allow developers to work at a higher level of abstraction and use component modelling early in the development process, which helps mitigate the complexities associated with MAS implementation [15], [16], [17], [18], [19], [20].

In MDE-based development processes, models are treated as first-class entities [21]. Engineers create models using various modelling languages that represent distinct parts of the system, providing a high level of abstraction. This abstraction enables engineers to concentrate on defining the system’s functionality rather than its implementation details. AOSE researchers have defined several agent metamodels [6], [7] to support modelling various aspects of MAS at appropriate abstraction levels. These metamodels are designed to capture agent characteristics such as plans, beliefs, goals, and interactions within MAS organizations.

To effectively implement MDE for MASs, a practical approach involves customizing Domain-Specific Modeling Languages (DSMLs) using integrated development environments (IDEs) that support modelling and code generation for the target system. Proposed MAS DSMLs (e.g., [22], [23], [24], [25],

[26], [27]) are based on the aforementioned agent metamodels and offer various abstract syntaxes. These DSMLs facilitate modelling both static and dynamic aspects of agent software from different perspectives within the MAS domain, including internal agent behaviour, interactions with other agents, and environmental entities.

Our research has focused on developing a platform-dependent modelling language that supports MASs, environment, and organization. We have developed a meta-model for JaCaMo that also leads to creating a syntax of a DSML for MAS development. To demonstrate its effectiveness, we have carried out qualitative evaluations.

III. THE SYNTAX OF THE LANGUAGE

This section introduces the metamodel of DSML4JaCaMo, representing the domain-specific language's abstract syntax. Generally, a metamodel outlines the system elements, their relationships, and cardinality constraints. It may also include attributes and operations for these elements. The metamodel is implemented using the EMF Ecore framework. Figure 1 illustrates the DSML4JaCaMo metamodel, describing the key meta-elements briefly. Detailed information regarding the JaCaMo components represented by the meta-elements can be reached by [1].

Beginning with Agent perspective, *Agent* defines the Jason BDI agent, which has *Plan*, *Belief*, *Rule*, *Goal*. A *Plan* element consists of *BodyTerm* and *Action*. As a *Plan* consists of a set of actions, a self-reference, namely *nextAction*, is created to model the action chain. An action can be an *InternalAction* and *ExternalAction*. *Message* element refers to agent communication. Lastly, *TriggeringEvent* meta-element is included as events are realized as consequences of beliefs or goals change in any Jason agent's mind.

In the Artifact dimension, a *Workspace* defines all contained artefacts and AbstractOperation E-Class, namely *AbsOperation* that the agents can utilize. *Port* element is used to link two artifacts within artifacts dimension. Since an *Artifact* element can have multiple *ObsProperty* elements that can be used in the operations such *LinkedOperation*, *InternalOperation*, *GuardOperation* and *Operation*.

On the Organisational perspective side, the Moise platform uses *XML* configuration comprising normative, structural and functional specifications. *NormativeSpecification* contains *Norm* elements which normalise the *Mission* elements. The *StructuralSpecification* has *Group* and *Role* EClasses in which subgroups and role extensions can be created. Moreover, *FormationConstraints* defines the features such as formation between roles and group scope. *Link* determines the link and their types among the roles. Lastly, *FunctionalSpecification* composes *Scheme* and *Scheme* creates *Mission* elements where each mission contains goals that agents achieve. In addition to *Scheme*, we also defined *OPlan* and *OGoal* EClasses, which allows us to model the layered goal and plan composition structure of Moise, which uses *XML* configuration. This composable layered structure via *OGoal* and *OPlan* elements are visualized in detail in section V.

The concrete syntax of a language encompasses the set of notations responsible for its presentation and construction. In the context of DSML specifications, the concrete syntax primarily facilitates mapping between meta-elements and their representations within instance models of the meta-model. Consequently, we developed a graphical concrete syntax that aligns the abstract syntax elements of DSML4BDI with their corresponding graphical notations. To achieve this, we leveraged the features of the Sirius¹ modelling environment. Sirius provides tools for creating a graphical editor from an Ecore metamodel and allows for the definition of specialized editors—including diagrams, tables, and trees—based on a viewpoint approach. This functionality enabled us to build the DSML4JaCaMo graphical modelling toolset within the Sirius environment for this study. Due to page limitations, we cannot provide the concrete system syntax as a table separately. However, you will see some examples in the case study graphs.

IV. CODE GENERATION: TRANSLATIONAL SEMANTICS

A comprehensive definition of a DSML cannot be achieved solely by specifying the notations and their representations. It also necessitates providing the semantics of language concepts, typically regarding the meanings of already established concepts. In this study, therefore, the metamodel elements are mapped to the concepts within the JaCaMo framework. This mapping between the metamodel and JaCaMo entities facilitates a series of model-to-text (M2T) transformations, constructing the DSML4JaCaMo's semantics within the JaCaMo framework. To this end, model-to-code transformation is used. Some excerpts of the generation rules are given and discussed below.

Listing 1: Excerpt from Acceleo rules for creating JaCaMo files

```

1 [template public generateElement(aMAS : MAS) ]
2 [comment @main/]
3 [for (ag : Agent | aMAS.agent)]
4 [file (ag.Name.concat('.asl'), false, 'UTF-8')]
5 .....
6 [for (p : Plan | ag.plan)]
7 [if (p.asBeliefAddition = 'true') + [else] + [if] [p.Name/]
  <- [p.hasContext.Expression/]
8 [if (p.hasBody.firstAction ->size()>0)]
9 [p.hasAction.thePlanSeq(p.hasAction, p.hasBody.
  firstAction) ->asOrderedSet().Expression /]
10 [/if]
11 [/for]
12 [/file]
13 [/for]
14 [for (wp : Workspace | aMAS.workspace)]
15 [for (art : Artifact | wp.artifact)]
16 [file (art.className.concat('.java'), false, 'UTF-8')]
17 import cartago.*;
18 public class [art.className/] extends Artifact {
19 .....
20 [for (Op: AbsOperation | art.operation)]
21 [if (Op.eClass().instanceTypeName.equalsIgnoreCase('
  dSML4JaCaMo.Operation'))]

```

¹Sirius modeling tool. Available at <https://eclipse.org/sirius/>, accessed May 2024


```

22 @OPERATION void [Op.className/]
23 {
24   ObsProperty prop = getObsProperty(" ");
25   prop.updateValue( );
26   signal (" tick ");
27 }
28 [/ if ]
29 .....
30 [/ if ]
31 [/ for ]
32 }
33 [/ file ]
34 [/ for ]
35 [/ for ]
36 [for (org : Organisation | aMAS.organisation)]
37 [ file (org.id.concat('.xml'), false, 'UTF-8')]
38 <?xml version="1.0" encoding="UTF-8"?>
39 <?xml-stylesheet href="http://moise.sourceforge.net/xml/
   os.xsl" type="text/xsl" ?>
40 <organisational - specification id="[org.id]/"
41 os-version="0.8" xmlns="http://moise.sourceforge.net/os"
42   xmlns:xsi="http://www.w3.org/2001/XMLSchema-
   instance"
43   xsi:schemaLocation="http://moise.sourceforge.net/
   os
44   http://moise.sourceforge.net/xml/os.xsd" >
45 [for (ss: StructuralSpecification | org.
   structuralSpecification )]
46 < structural - specification >
47 .....
48 </ structural - specification >
49 [/ for ]
50 < functional - specification >
51 [for (fs: FunctionalSpecification | org.
   functionalSpecification )]
52 [for (sch: Schemel fs.scheme)]
53 <scheme id="[sch.id]/" >
54 [for (Og: OGoal | sch.SchemeOgoal )]
55 [if (Og.isRootGoal)]
56 <goal id="[Og.Name/]" ttf="5 seconds">
57 [for (Op: OPlan | Og.OGoalToOPlan )]
58 <plan operator="[if (Op.Parallel)] parallel [else]sequence
   [if]" >
59 [for (FOg: OGoal | Op.FirstOgoal )]
60 <goal id="[FOg.Name/]" ds="" >
61 .....
62 </goal>
63 [/ if ]
64 .....
65 [/ for ]
66 [/ for ]
67 </ functional - specification >
68 [for (ns: NormativeSpecification | org.
   normativeSpecification )]
69 [for (norm: Norm | ns.norm )]
70 <normative- specification >
71 .....
72 </normative- specification >
73 [/ for ]
74 [/ for ]
75 [/ file ]
76 [/ for ]
77 [/ template ]

```

Within the JaCaMo framework, each agent is represented by an ASL file, which includes code written in the AgentSpeak

language, designed to outline the internal structure of the agent according to the BDI architecture. Fundamentally, an AgentSpeak agent is characterized by a set of beliefs, rules, and plans. Beliefs denote the initial knowledge possessed by the agent. Rules are logical expressions or mathematical equations that guide the agent's reasoning. Plans consist of the actions and/or subgoals the agent employs to achieve its current objectives. Each plan within an AgentSpeak agent comprises a triggering event, a context, and a body element. The triggering event indicates the circumstances under which the plan is appropriate. The context determines the plan's applicability based on the agent's beliefs. The body is a sequence of basic actions and/or subgoals the agent will execute.

Listing 1 includes an excerpt from the Acceleo rules to generate all Agent ASL, Artifact Java and Organisation XML files. Initially, in lines 3 to 16, an ASL file is created for each agent. Lines 6 to 11 are responsible for generating each agent's plans, plans' contexts and the actions in these plans, including their triggering events and goals.

Lines 14 to 35 are responsible for writing the necessary Java code for the artefacts of each workspace in the environment. Each Artefact Java file includes Cartago library and Artefact naming, which extends the Artifact class. The generator then synthesis operations that can be *Operation*, *Linked Operation*, *Guard Operation* and *Internal Operation*. The generator checks the EClass of these elements to generate the necessary type. Lastly, observable properties are also generated within the corresponding operation types.

The remaining lines, defined between 36 and 76, are responsible for creating the organization file, including the agents' roles in the system, their responsibilities, and the organizational goals and plans, aforementioned as *OGoal* and *OPlan* elements. The code generation for organisational dimension also considers specifications which are *Functional*, *Structural* and *Normative*. Specifically, an excerpt from *StructuralSpecification* is defined lines between 45-49, and *FunctionalSpecification* is defined between 50 and 67, and lines 68-77 scopes the *NormativeSpecification*. As *StructuralSpecification* creates goal composition and mapping relations such as *OGoalToOPlan* and *FirstOGoal* is designed to preserve this structure at the modelling level as well. This structure is exemplified in Figure 2. Due to size and space limitations, a brief description of the transformation rules is mentioned.

V. CASE STUDIES & EVALUATION

Two case studies are planned and implemented with DSML4JaCaMo to validate the system and assess its code-generating capability. The case study models are displayed.

Figure 4 shows how the DSML builds the code artifacts based on the intended models and how the delta codes are incorporated to create the final code.

Our modelling language is evaluated by examining its code-generation capabilities. This involves comparing the lines of generated code with the final code, which includes the additional delta code.

A. Case 1: Writing Paper

In this case study, a set of agents coordinate to mimic the process of writing a paper. This is achieved with the help of assistant agents following a structured plan and having different roles and missions. There are roles such as writer, editor and manager. The writing goal is divided into multiple goals achieved via agents’ plans. In addition, an agent uses a checkList artifact to mark the finalisation of paper writing. Figure 2 depicts the organisational structure, i.e., the scheme of the case study.

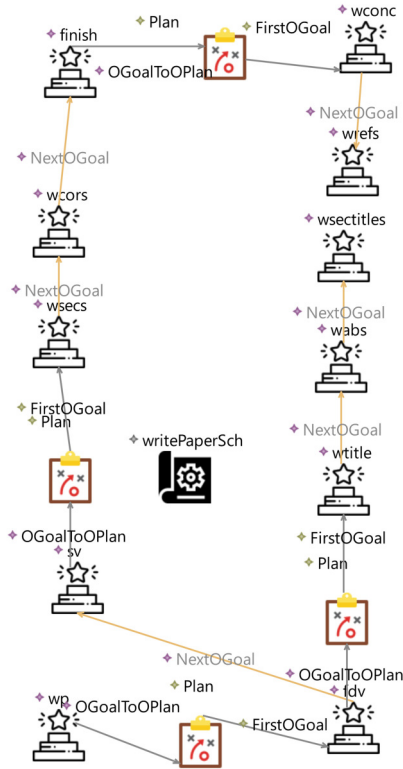


Fig. 2: Scheme design based on OGoals and OPlans composition.

B. Case 2: Harvest Process for Pizza

In this case study, agents coordinate to mimic the pizza process, from wheat harvesting to dough preparation to having a pizza. The pizza-making goal is broken down into multiple goals. In addition, an agent uses an oven artifact to mimic the cooking process based on/off control. Table I shows the code generation performance for Jason, Cartago and Moise platforms based on lines of code. In each scenario, 3 agents for Jason, 1 Scheme for Moise and 1 Artifact for Cartago were used. Figure 3 illustrates these three agents connected to a Cartago Workspace and joined the organisation, namely *Org 1*. It uses a scheme called *harvestForPizza*, where three norms and three roles are defined. In addition, Figure 4 shows the Agent viewpoint that contains the Plans and their corresponding contexts.

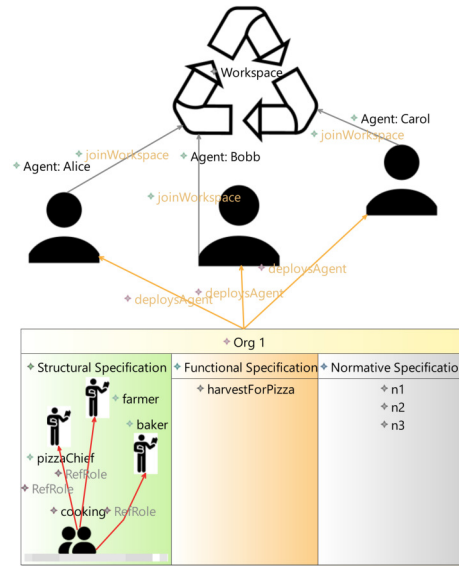


Fig. 3: The MAS viewpoint of the harvest case study.

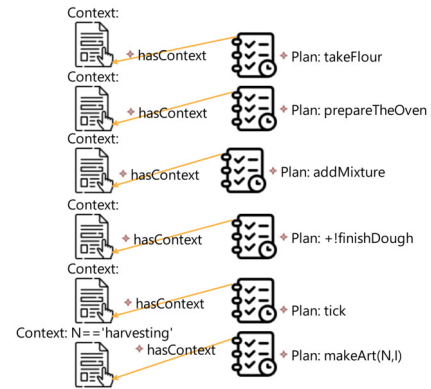


Fig. 4: Agent viewpoint of the harvest case study.

Eventually, based on model-to-text transformation, we assessed the code generation performance using two case studies. Although this facilitates the system’s development, another MDE approach, the model-to-model transformation, is also required to verify the JaCaMo. Specifically, the designed Jason agents and Moise organisation need static analyses before the system’s deployment. In this regard, Coloured Petri-nets (CPN) are a suitable paradigm to transform the JaCaMo model to CPN representations to achieve analyses [5], [28] using a non-deterministic domain. Another approach is that Moise’s

TABLE I: Evaluation results for two case studies based on lines of code.

	Agent 1	Agent 2	Agent 3	Agent Total	1 Scheme	1 Artifact
Case1 Normal	6	35	3	42	77	11
Case 1 Generated	6	28	3	37	60	9
Case 2 Normal	12	38	3	53	74	23
Case 2 Generated	12	24	3	39	44	21

Scheme design needs a proper formalism for representation. For this purpose, Statecharts can be used as there is quite a similarity between deterministic OGoal and OPlan structures. In the next section the paper is concluded.

VI. CONCLUSION

This study presents a DSML called DSML4JaCaMo, detailing its abstract and concrete syntaxes with translational semantics. The DSML is evaluated through two case studies, showing that 76% of the total system code is generated automatically. Our study offers an abstraction to simplify complexity and alleviate difficulties. We use a set of graphical notations and domain constraints to create a graphical editor for the DSML. The effectiveness of the DSML's generation capability is assessed through two case studies. In future work, we aim to provide a more comprehensive evaluation of the language and its tool by following the approach for the multi-case systematic evaluation of MAS DSMLS introduced in [19]. That evaluation will provide the quantitative measurement of the language's usability as well as the assessment of the language features according to a series of well-defined quality characteristics for MAS development.

REFERENCES

- [1] O. Boissier, R. H. Bordini, J. Hubner, and A. Ricci, *Multi-agent oriented programming: programming multi-agent systems using JaCaMo*. MIT Press, 2020. ISBN 9780262044578
- [2] P. Leitao, S. Karnouskos, L. Ribeiro, J. Lee, T. Strasser, and A. W. Colombo, "Smart agents in industrial cyber-physical systems," *Proceedings of the IEEE*, vol. 104, no. 5, pp. 1086–1101, 2016. doi: 10.1109/JPROC.2016.2521931
- [3] B. Karaduman, I. David, and M. Challenger, "Modeling the engineering process of an agent-based production system: An exemplar study," in *2021 ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. IEEE, 2021. doi: 10.1109/MODELS-C53483.2021.00051 pp. 296–305.
- [4] M. Brambilla, J. Cabot, and M. Wimmer, "Model-driven software engineering in practice," *Synthesis lectures on software engineering*, vol. 3, no. 1, 2017. doi: <https://doi.org/10.1007/978-3-031-02549-5>
- [5] B. Karaduman, B. T. Tezel, and M. Challenger, "Towards static analysis of bdi agents on cps using petri nets and model-driven engineering," in *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer (in press), 2024.
- [6] C. Hahn, C. Madrigal-Mora, and K. Fischer, "A platform-independent metamodel for multiagent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 18, no. 2, pp. 239–266, 2009.
- [7] B. T. Tezel, M. Challenger, and G. Kardas, "A metamodel for Jason BDI agents," in *5th Symposium on Languages, Applications and Technologies (SLATE'16)*, vol. 51, 2016. doi: 10.4230/OASlcs.SLATE.2016.8 pp. 8:1–8:9.
- [8] O. Boissier, R. H. Bordini, J. F. Hübner, and A. Ricci, "Dimensions in programming multi-agent systems," *The Knowledge Engineering Review*, vol. 34, p. e2, 2019. doi: 10.1017/S026988891800005X
- [9] O. Shehory and A. Sturm, *Agent-Oriented Software Engineering: Reflections on Architectures, Methodologies, Languages, and Frameworks*. Springer-Verlag Berlin Heidelberg, 2014.
- [10] B. Lelandais, M.-P. Oudot, and B. Combemale, "Applying model-driven engineering to high-performance computing: Experience report, lessons learned, and remaining challenges," *Journal of Computer Languages*, vol. 55, p. 100919, 2019. doi: <https://doi.org/10.1016/j.cola.2019.100919>
- [11] A. Bucchiarone, J. Cabot, R. F. Paige, and A. Pierantonio, "Grand challenges in model-driven engineering: an analysis of the state of the research," *Software and Systems Modeling*, vol. 19, no. 1, pp. 5–13, 2020. doi: <https://doi.org/10.1007/s10270-019-00773-6>
- [12] C. Verbruggen and M. Snoeck, "Model-driven engineering: A state of affairs and research agenda," *Enterprise, business-process and information systems modeling*, 2021. doi: 10.1007/978-3-030-79186-5_22
- [13] E. de Araújo Silva, E. Valentin, J. R. H. Carvalho, and R. da Silva Barreto, "A survey of model driven engineering in robotics," *Journal of Computer Languages*, vol. 62, 2021. doi: <https://doi.org/10.1016/j.cola.2020.101021>
- [14] D. Di Ruscio, D. Kolovos, J. de Lara, A. Pierantonio, M. Tisi, and M. Wimmer, "Low-code development and model-driven engineering: Two sides of the same coin?" *Software and Systems Modeling*, vol. 21, no. 2, 2022. doi: <https://doi.org/10.1007/s10270-021-00970-2>
- [15] Y. E. Cakmaz, O. F. Alaca, C. Durmaz, B. Akdal, B. Tezel, M. Challenger, and G. Kardas, "Engineering a bdi agent-based semantic e-barter system," in *2017 International Conference on Computer Science and Engineering (UBMK)*. IEEE, 2017. doi: 10.1109/UBMK.2017.8093474 pp. 1072–1077.
- [16] B. T. Tezel, M. Challenger, and G. Kardas, "Dsm14bdi: A modeling tool for bdi agent development," in *12th turkish national software engineering symposium (uyms 2018)*, 2018, pp. 1–8.
- [17] V. Mascardi, D. Weyns, A. Ricci, C. B. Earle, A. Casals, M. Challenger, A. Chopra, A. Ciorca, L. A. Dennis, Á. F. Díaz *et al.*, "Engineering multi-agent systems: State of affairs and the road ahead," *ACM SIGSOFT Software Engineering Notes*, vol. 44, no. 1, pp. 18–28, 2019. doi: <https://doi.org/10.1145/3310013.3322175>
- [18] M. Challenger, B. T. Tezel, V. Amaral, M. Goulao, and G. Kardas, "Agent-based cyber-physical system development with sea_ml++," in *Multi-Paradigm Modelling Approaches for Cyber-Physical Systems*. Elsevier, 2021, pp. 195–219.
- [19] O. F. Alaca, B. T. Tezel, M. Challenger, M. Goulão, V. Amaral, and G. Kardas, "Agentdsm-eval: A framework for the evaluation of domain-specific modeling languages for multi-agent systems," *Computer Standards & Interfaces*, vol. 76, p. 103513, 2021. doi: <https://doi.org/10.1016/j.csi.2021.103513>
- [20] B. Karaduman, B. T. Tezel, and M. Challenger, "Rational software agents with the bdi reasoning model for cyber-physical systems," *Engineering Applications of Artificial Intelligence*, vol. 123, p. 106478, 2023. doi: <https://doi.org/10.1016/j.engappai.2023.106478>
- [21] G. Kardas, F. Ciccozzi, and L. Iovino, "Introduction to the special issue on methods, tools and languages for model-driven engineering and low-code development," *Journal of Computer Languages*, vol. 74, 2023. doi: <https://doi.org/10.1016/j.cola.2022.101190>
- [22] M. Challenger, S. Demirkol, S. Getir, M. Memrik, G. Kardas, and T. Kosar, "On the use of a domain-specific modeling language in the development of multiagent systems," *Engineering Applications of Artificial Intelligence*, vol. 28, pp. 111–141, 2014. doi: <https://doi.org/10.1016/j.engappai.2013.11.012>
- [23] E. J. T. Gonçalves, M. I. Cortés, G. A. L. Campos, Y. S. Lopes, E. S. Freire, V. T. da Silva, K. S. F. de Oliveira, and M. A. de Oliveira, "Mas-ml 2.0: Supporting the modelling of multi-agent systems with different agent architectures," *Journal of Systems and Software*, vol. 108, pp. 77–109, 2015. doi: <https://doi.org/10.1016/j.jss.2015.06.008>
- [24] F. Bergenti, E. Iotti, S. Monica, and A. Poggi, "Agent-oriented model-driven development for JADE with the JADEL programming language," *Computer Languages, Systems & Structures*, vol. 50, pp. 142–158, 2017. doi: 10.1016/j.cl.2017.06.001
- [25] G. Kardas, B. T. Tezel, and M. Challenger, "Domain-specific modelling language for belief-desire-intention software agents," *IET Software*, vol. 12, no. 4, pp. 356–364, 2018. doi: <https://doi.org/10.1049/iet-sen.2017.0094>
- [26] D. Sredojević, M. Vidaković, and M. Ivanović, "Alas: agent-oriented domain-specific language for the development of intelligent distributed non-axiomatic reasoning agents," *Enterprise Information Systems*, vol. 12, no. 8-9, pp. 1058–1082, 2018. doi: <https://doi.org/10.1080/17517575.2018.1482567>
- [27] A. Siabdelhadi, A. Chadli, H. Cherroun, A. Ouared, and H. Sahraoui, "Motrans-bdi: Leveraging the beliefs-desires-intentions agent architecture for collaborative model transformation by example," *Journal of Computer Languages*, vol. 74, p. 101174, 2023. doi: <https://doi.org/10.1016/j.cola.2022.101174>
- [28] B. Karaduman, M. Challenger, R. Eslampanah, J. Denil, and H. Vangheluwe, "Analyzing WSN-based IoT Systems using MDE Techniques and Petri-net Models," in *4th International Workshop on Model-Driven Engineering for the Internet-of-Things (MDE4IoT), Co-Located With Software Technologies: Applications and Foundations (STAF 2020), Virtual Event, Norway*, 2020, pp. 35–46.

Towards understanding animal welfare by observing collective flock behaviors via AI-powered Analytics

Savvas Karatsiolis^{1,*}, Pieris Panagi²,
Vassilis Vassiliades³, Andreas Kamilaris^{4,5}

¹0000-0002-4034-7709

²0009-0007-4027-1488

³0000-0002-1336-5629

⁴0000-0002-8484-4256

^{1,2,3,4}Cyens Centre of Excellence

1 Plateia Dimarchou Lellou, Nicosia 1016, Cyprus

⁵Department of Computer Science, University of Twente,
7522 NB Enschede, The Netherlands

Email: {s.karatsiolis, p.panagi, v.vassiliades, a.kamilaris}@cyens.org.cy

* Corresponding author

Nicolas Nicolaou¹, Efstathios Stavrakis²

¹0000-0001-7540-784X

²0000-0002-9213-7690

Algolysis Ltd,

Archiepiskopou Makariou III 200, Lakatamia 2311,
Nicosia, Cyprus

Email: {nicolas, stathis}@algolysis.com

Abstract—Animal farming has undergone significant transformation and evolved from small-scale businesses to large-scale commercial ventures. While maximizing productivity and profitability has always been a major concern in animal farming, during recent years there has been an increasing rise of concern regarding the welfare of the animals. In this context, the integration of artificial intelligence (AI) technologies offers immense potential for monitoring the well-being of chickens on farms and optimizing revenue streams simultaneously. Several works have integrated AI methodologies into everyday animal farming activities. Still, very few (if any) have proposed efficient and practical solutions that may facilitate farm owners in making impactful decisions regarding their business profitability and the welfare of the animals. In this direction, we propose a non-invasive chicken farm monitoring system that relies on onfield sound and video recordings integrated with sensory data acquired from the farm. The system consists of hardware that handles data acquisition and storage, a sensory data collection system and audio/video processing AI models. The last component of the system will be an inference engine that analyzes the collected data and infers useful facts about the flock’s welfare and even psychological state.

I. INTRODUCTION

Ensuring the well-being of chickens on a farm is of paramount importance for ethical, environmental, and economic reasons. The welfare of chickens directly impacts the quality of the final product, e.g., the number and the quality of the daily produced eggs or the meat produced. Chicken flocks living in a clean environment, being fed adequately, having space to roam and being stress-free are less susceptible to diseases and premature death [1]. In this work, we explore monitoring the well-being of the chickens on the farm with technology to understand whether the flock is under stress and identify the source(s) of stress. This enables the farm management to deal with potential stress-inducing factors, keep the flock healthy and prevent catastrophic consequences. AI-driven monitoring systems can play a crucial role in this regard by continuously assessing implicit health or behavior

cues like the flock’s clucking or its daily motion index [12] along with environmental conditions. By leveraging machine learning algorithms, such systems can identify potential health issues or disease outbreaks, enabling proactive measures to be taken. Good, non-invasive information sources for inferring the well-being of the animals on a chicken farm are audio and video recordings of the chicken combined with the data acquired by sensors measuring ambient pollutants and the weather conditions on the farm. Chickens are especially vocal birds and tend to express their psychological and physical state through clucking to each other which enables the monitoring of their welfare and the detection of stressful conditions. Finally, the health and welfare of the birds on a farm rely on the conditions they live in. For example, high concentrations of ammonia on a farm [11], extreme (low or high) temperatures and humidity affect the health of the chickens and should be monitored carefully.

II. RELATED WORK

Many research works on AI-empowered chicken farm monitoring use computer vision to detect individual animals and track their motion on the farm. For example, [5] proposes applying optical flow in video recordings of chickens for identifying early signs of infection by the pathogen *Campylobacter*. Along similar lines, the authors of [9] identify early warnings of footpad dermatitis and hockburn in broiler chicken flocks.

In the era of Deep Learning, several research works suggest using popular Neural Network architectures to conduct individual chicken detection from video footage collected in the field. For example, in [3] the authors use Faster R-CNN [15] architectures to develop object detection and instance segmentation models that operate on edge devices installed on the farm. They propose combining these models with the monitoring of environmental parameters for early disease detection which is the subject of future work in those papers. A similar work [19] proposes the use of the YOLOv5 architecture

[17] to detect cage-free chickens on the litter floor. Besides the provision of the technology for non-invasive inference based on video recordings, several researchers propose methodologies in the context of Precision Livestock Farming (PLF) which involves wireless and Radio Frequency Identification (RFID) sensors on the chickens. For example, the authors of [4] detect the movements of each chicken with an RFID system and classify them into active, normal, or sick, claiming they can detect sick chickens at early sickness stages before the whole flock is affected. Another interesting approach by [7] involves automated monitoring and quantification of feeding and nesting behaviors of individual hens.

Finally, animals' vocalizations can be exploited as they contain a wealth of biological information (e.g., reflecting their social interactions, communicating alarm signals and containing cues about their psychological state). In this context, animals' vocalization may be used as a welfare indicator [10]. The clucking of chickens is a form of communication within the flock; it conveys messages between the birds and can be used as a warning signal or a means of conveying messages of discomfort, stress, satisfaction and expressing social interactions among the chicken.

Our paper differs from related work in that we employ advanced AI techniques to combine multiple modalities of sensory observations (i.e., audio and video recordings) from the farm environment to enable the production of analytics.

III. TOOLS AND METHODOLOGY

A. Hardware

We employ audio and video recordings of the flock to conduct inference on the condition of the animals on the farm. For data acquisition, we set up a network of sensory devices each equipped with a microphone, a camera, a speaker¹ and a local storage device (Solid State Drive-SSD) for temporarily storing collected data. The heart of the system is a centralized device that synchronizes the acquisition conducted by the sensory devices over the network via Application Programming Interface (API) calls. The centralized device also acts as a Network Access Server (NAS) and Processing Engine (PE) that runs the audio Neural Network (NN) and the motion detection algorithm on the data (audio and video) captured by the sensory devices and sent to the NAS. We designed the hardware so that important configuration parameters (e.g., acquisition interval and duration, external stimuli sound, network configuration like sensors' IP addresses, etc.) are configurable on the centralized device. The centralized device will be referred to as the Synchronization and Processing Engine (SPE). The SPE performs the following:

- Implements API calls for conducting synchronization.
- Records status/error logging from sensors' communication.
- Hosts the NAS service.
- Runs inference on the audio NN.

¹The speakers will be used in the future for producing short sounds to assess the flock's response to external stimuli




Fig. 1. The communication flow between the SPE and the sensors in the monitoring network. The SPE handles the administrative tasks and overall coordination. Every message exchange is carried out via API calls.

- Runs the motion detection algorithm.

The sensor devices capture audio/video streams in response to the synchronization messages sent by the SPE (the synchronization messages also instruct the sensors of the duration of the imminent video/audio acquisitions). For all devices in the monitoring network, we use Raspberry Pi (RPI) [13] modules because they are flexible, support audio/video acquisition and provide a fair amount of processing power. Due to the more demanding operations required by the SPE, it is built around RPi model 5 while the sensor modules are built around RPi model 4. The operational flow of the data acquisition hardware is shown in Figure 1. To efficiently handle the role of the NAS, the SPE is equipped with a 16TB Hard Disk Drive (HDD) to store several weeks' worth of data. It also supports a hot swap operation i.e., another HDD can replace the working HDD while the system is operational. This enables the operators of the monitoring system to transport the data to other premises and further process it. We will soon publish the audio/video datasets for general public use.

B. Deployment

We built a small-scale monitoring system comprising one SPE and two sensor modules and installed it on a chicken farm located in a rural area in the district of Nicosia, Cyprus. All monitoring system devices are placed in protective enclosures (fabricated with 3D printers) to deal with the harsh environmental conditions on the farm, i.e., large concentrations of sand and particles in the air, high humidity and extreme temperatures. Figure 2 shows images of the installation on site.

We configured the system to acquire video and audio recordings of 58 seconds per minute and allow 2 seconds per minute



Fig. 2. Photos of the monitoring system.

for system operations. We record audio data throughout the day (24 hours) and video for 12 hours per day only (from 06:00 in the morning to 6:00 p.m.) because our cameras operate in the visible light spectrum. Although chickens are mostly quiet at night, we keep recording audio during night hours to detect potential nocturnal predators' attacks. To train/develop our audio and video models, we gathered data for 50 days which translates to a dataset comprising 75K audio and 37.5K video datapoints.

C. Audio Anomaly Detection Model

We acknowledge the difficulty of creating an accurate dataset containing chicken vocalizations annotated with their psychological state, e.g., satisfied, calm, stressed, or panicked. Developing such a dataset requires the contribution of several experts and great effort. Alternatively, we use unsupervised learning for training the model which is much cheaper and faster and utilizes the acquired data more efficiently. Thus, the data is used to create a benchmark for the vocalization of chickens on a certain farm. Then, we exploit this benchmarked vocalization database to infer anomalies in future chicken vocalizations.

1) *Data pre-processing*: The collected audio streams are transformed into Mel spectrograms [14] and then processed by a Convolutional Denoising Autoencoder (Conv-DAE) [18]. Mel spectrograms apply a frequency-domain filter bank to audio signals that are windowed in time. We use Mel spectrograms because they offer a more perceptually relevant representation of audio signals being aligned with human sound perception. Essentially, the audio streams are transformed into frequency domain filter banks that describe the sound signal in terms of its frequency content. Some examples of the Mel spectrograms created from audio files acquired on the farm are shown in Figure 3. The Mel spectrograms provide a signature of the processed audio that reflects the psychological state of the chicken and thus enable the efficient learning of audio features in chicken clucking or any other sound made by the birds. To learn these audio features, we use unsupervised learning and thus no annotations are required. Specifically, we apply a Conv-DAE that learns to reconstruct the Mel spectrograms from their noisy versions.

Given a Mel spectrogram x sampled from a dataset X , the model accepts an input $x' = x + \mathcal{N}(\mu, \sigma^2)$, with \mathcal{N} being a Gaussian noise process, and outputs \hat{x} . The model learns to minimize the Mean Squared Loss (MSE) between




Fig. 3. The Mel spectrograms of 2-second chunks of audio streams collected on the farm. Each spectrogram shows a different frequency content in the audio files. We use these frequency representations to learn the features of the audio and infer the psychological condition of the chicken.

the reconstructed spectrogram and the original spectrogram. Concretely, the optimization objective function is $\min L = \min \sum_i \frac{1}{2}(\hat{x}_i - x_i)^2, i \in [0, |X|]$. In our experiments, we use $\mu = 0, \sigma^2 = 0.25$ and compute the Mel Spectrograms on 2-second chunks of audio data.

Since each sensor module records an audio sequence, we concatenate all individual Mel spectrograms obtained from the sensor modules into one data structure that has multiple channels. The number of channels equals the number of obtained spectrograms (the number of the system's sensor modules). This technique allows using a single NN that accounts for all sound data recorded by the sensors. The main advantage of processing all sound streams concurrently instead of processing them separately is the inherent ability to combine sounds from different locations on the farm that represent the same event, i.e., the model gets recordings of the same observation from different standpoints which, in certain cases, may provide richer information. For example, a loud sound made by a chicken may overtake the signal recorded by a certain nearby microphone preventing the device from representing any other equally important sounds sourced from different locations inside the farm.

2) *Modeling*: By learning to restore the original spectrograms, the model learns the features of the problem domain and thus becomes capable of identifying the peculiarities of the data. In other words, the model learns the manifold of the data and distills the low-level audio characteristics that comprise the data [2]. To infer the flock's psychological state we observe the relative location of audio samples mapped to the benchmarked feature space on the data manifold. Since different positions on the manifold reflect different acoustic characteristics, the psychological state of the chicken(s) expressed by certain vocal characteristics can be inferred based on the mapping of the data. To create the Mel spectrograms, we first split each 58-second audio recording into 2-second chunks. We chose a 2-second audio interval because it is a good fit for providing an audio signature: it is adequate for providing sufficient vocal information while not being big enough to cause severe audio frequency shifts that may harm the processing. We process each chunk with 80 spectral banks and fast Fourier transforms of size 2220 to produce spectrograms of size 80×80.

The Conv-DAE is comprised of an encoder-decoder architecture. The encoder maps the input (noisy Mel spectrograms) to a latent space and the decoder reconstructs the input. For




Fig. 4. The architecture of the Convolutional Autoencoder used to learn the features of the audio acquired on the farm. By learning to reconstruct noisy Mel spectrograms of recorded audio, the model builds a knowledge of the underlying elements of chickens' vocal cues.

the encoder, we use a Residual Network [6] and specifically the ResNet18 variant and for the decoder a sequence of Convolutional and Up-sampling layers. The model contains 26M trainable parameters and is shown in Figure 4.

3) *Training*: The encoder compresses the 80X80 spectrograms into a 512-D representation and the decoder reconstructs the original spectrogram by processing the 512-D vector. To train the model, we split the acquired audio dataset into a training and a testing set (80%-20% respectively) and the latter is used for assessing the model. The model training achieves a mean absolute error (MSE) of 0.126 while the test set error is 0.141. To evaluate the quality of the embeddings produced by the encoder we reduce the dimensionality of the test set embeddings from 512-D to 3-D with Principal Components Analysis (PCA) by projecting the data onto a much smaller subspace. This allows the visualization of the embedding space for evaluation purposes. We further identify the top 100 points with the highest reconstruction error and perceive them as anomalous cases: these are the data points that the model cannot adequately reconstruct which means that they are either outside the high probability regions of the distribution or they are under-represented in the dataset.

4) *Visualizations*: We exploit the low-dimensional (3-D) representations of the test set calculated by PCA to link the audio semantics of the data points with their distribution in space (left image in Figure 5). Interestingly, the data points in the test set are distributed in the 3-D latent space in such a way that data points with similar semantics are feature-mapped close to each other. For example, most of the audio data points that contain very soft clucking are mapped close to each other at a certain region of the latent space. Likewise, most of the audio data points that contain the clucking of rather stressed chicken are mapped close to each other in a certain region of the feature space. Most importantly, most of the anomalous points (data points with the highest reconstruction error) contain sounds of panicked birds that make distinct sounds of despair. Still, this 3-D feature mapping is imperfect because it does not distinguish between the data points in a definite and clear manner mostly because of the information loss during the dimensionality reduction and the confusing mix-up of chicken sounds and environmental noises (especially noises from machinery and feeding/watering equipment).

5) *Clustering*: We further cluster the 3-D embeddings with the k -means algorithm [8] into 5 regions. The choice of using 5 regions lies with the way the embeddings are spread onto the feature space. Figure 5 (right image) shows the 3-D




Fig. 5. Left: The 3-D embeddings of the test dataset after applying dimensionality reduction for visualization purposes. The points shown in red are perceived as anomalous points because they have high reconstruction errors. Right: The 3-D embeddings are clustered with the k -means algorithm. The resulting clusters shown with various colors represent different sound semantics.

embeddings of the test set and the clustering obtained with the k -means algorithm.

The clusters calculated by the k -means contain semantically different sounds:

- Red cluster: Low-intensity sounds (flock resting and being very calm).
- Blue cluster: Normal soft clucking.
- Yellow cluster: Calm clucking and ambient noises (like food-delivery-machinery).
- Black cluster: Flock noises ranging from clucking of medium intensity to extremely loud flock sounds (panic sounds).
- Green cluster: Very soft clucking and ambient noise (mainly fans blowing air in the farm to cool down the flock)

Most importantly, we observe that the anomalous points (the ones with the highest reconstruction loss, shown in red color in the left image of Figure 5) are located at the extremities of the black cluster (medium to extreme noises). We provide a video that demonstrates this analysis with sound to show the different sound semantics of the various clusters at <https://sworld.cyens.org.cy>.

D. Motion Detection

Similar to the case of audio recordings, video can also be used to detect whether the chickens in a farm are calm or under stress. The primary indicator of the stress level in video recordings is the chickens' motion: stressed chickens make rapid movements, wander around loudly and become very jerky. We detect chickens' motion with an algorithm based on background removal. Background removal or subtraction is commonly used to segment moving parts from static scenes (background and foreground). The motion is detected by subtracting the current frame of the video from a reference static background calculated by a background modeling technique which is continuously updated. One of the most popular background subtraction algorithms is the Mixture of Gaussians (MOG). According to MOG [16], for each background pixel,



Fig. 6. Left: A single frame from a video recording from the farm. Right: The motion mask produced by MOG2.

a mixture of Gaussian distributions and a weighting parameter are utilized to "save the lifetime of pixels in the scene". Pixels with a long lifetime are interpreted as background pixels while the rest are characterized as foreground or pixels belonging to objects that move around. An improved version of MOG is called MOG2 [20] and determines the appropriate number of distributions for the modeling.

We apply MOG2 on the video streams we acquired from the sensor modules, to obtain the motion masks (images showing where motion is detected) of the frames comprising the videos. We use a relatively high video frame rate (30 frames per second) to facilitate MOG2 in detecting subtle chicken movements. Figure 6 shows the output of the motion detection process applied on a single random frame.

The motion masks produced by the MOG2 algorithm are binary images indicating whether there is motion at a certain pixel of the input frame. To generate a meaningful metric that reflects chickens' motion, we aggregate the masks by adding them together and thus compute a single number representing the magnitude of the flock's motion during the interval of the processed recording: a higher value means a more active flock. A video demonstration of the motion detection of the chicken can be found at <https://sworld.cyens.org.cy>.

Quantifying the flock's motion makes it possible to set an activity threshold which, when exceeded, could indicate that the flock is reacting to a threat or stressful condition. Symmetrically, another activity threshold could be set which, when it is not reached, could indicate that the flock is indolent and may suffer from a disease or being stressed by abnormal environmental conditions or underfeeding. Furthermore, the flock motion could be used to calculate motion statistics at different intervals (hourly, daily, weekly, etc.) and thus infer the activity level and the stress of the chicken on the farm from different perspectives.

IV. DISCUSSION-CONCLUSIONS

This paper introduced two methods for chicken flock monitoring, which could potentially lead to animal welfare indicators. Both the audio-based anomaly detection and the video-based motion detection methods directly provide the means to generate alerts to the farm personnel regarding the status of the flock. Besides the real-time detection of unpleasant situations, the combination of multi-modal data (audio streams, video streams and sensory data) into a unified system can be utilized

to assess the welfare of the chicken. By combining the data originating from multiple modalities and collaborating with experts in the field, we expect to get reliable indicators that can then be exploited by farm owners to manage their farms better, improving the welfare of their flock, increasing revenue and preventing catastrophic events.

Analyzing the outputs of our proposed system would provide powerful analytics that humans cannot easily observe while managing the farm. To our knowledge, those analytics are not available by any commercial system to date. Some examples of the analytics we aspire to provide to the farm managers/owners in the future are the following:

- The average flock motion this week was 30% less than the average daily flock motion during last week. *This may be a sign of underfeeding, extreme temperature, or environmental pollution like high ammonia.*
- The average flock motion during the last three days was 70% less than last month's average. *This may indicate a serious condition like illness. Immediate action needs to be taken.*
- Average flock motion during feeding time is 140% higher than usual. *Maybe, this is a sign of underfeeding.*

These examples reflect our reasoning that the outputs of our system need to be translated into domain-specific analytics by experts in poultry farming.

Funding: The work was funded by the European Union Recovery and Resilience Facility of the NextGenerationEU instrument, through the Research and Innovation Foundation (CODEVELOP-ICT-HEALTH/0322/0061), as well as the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No. 739578, and the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy.

REFERENCES

- [1] Haitham G. Abo-Al-Ela, Seham El-Kassas, Karima El-Naggar, Safaa E. Abdo, Ali Raza Jahejo, and Rasha A. Al Wakeel. "Stress and immunity in poultry: light management and nanotechnology as effective immune enhancers to fight stress". In: *Cell Stress and Chaperones* 26.3 (2021), pp. 457–472. ISSN: 1355-8145. DOI: <https://doi.org/10.1007/s12192-021-01204-6>. URL: <https://www.sciencedirect.com/science/article/pii/S1355814523005187>.
- [2] Guillaume Alain and Yoshua Bengio. "What regularized auto-encoders learn from the data-generating distribution". In: *Journal of Machine Learning Research* 15.1 (2014), pp. 3563–3593. DOI: 10.5555/2627435.2750359. URL: <https://dl.acm.org/doi/10.5555/2627435.2750359>.

- [3] Stevan Cakic, Tomo Popovic, Srdjan Krco, Daliborka Nedic, Dejan Babic, and Ivan Jovovic. "Developing Edge AI Computer Vision for Smart Poultry Farms Using Deep Learning and HPC". In: *Sensors* 23.6 (2023). ISSN: 1424-8220. DOI: 10.3390/s23063002. URL: <https://www.mdpi.com/1424-8220/23/6/3002>.
- [4] Zhang Feiyang, Hu Yueming, Chen Liancheng, Guo Lihong, Duan Wenjie, and Wang Lu. "Monitoring behavior of poultry based on RFID radio frequency network". In: *International Journal of Agricultural and Biological Engineering* 9.6 (2016), pp. 139–147.
- [5] Colles Frances et al. "Monitoring chicken flock behaviour provides early warning of infection by human pathogen". In: *Royal Society B: Biological Sciences* 283 (1822 2016).
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Deep Residual Learning for Image Recognition". In: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. IEEE Computer Society, 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90. URL: <https://doi.org/10.1109/CVPR.2016.90>.
- [7] Lihua Li, Yang Zhao, Jofran Oliveira, Wilco Verhoijssen, Kai Liu, and Hongwei Xin. "A UHF RFID system for studying individual feeding and nesting behaviors of group-housed laying hens". In: *Transactions of the ASABE* 60.4 (2017), pp. 1337–1347.
- [8] Stuart P. Lloyd. "Least squares quantization in PCM". In: *IEEE Trans. Inf. Theory* 28.2 (1982), pp. 129–136. DOI: 10.1109/TIT.1982.1056489. URL: <https://doi.org/10.1109/TIT.1982.1056489>.
- [9] Dawkins M.S., Roberts S.J., Cain R.J., Nickson T., and Donnelly C. A. "Early warning of footpad dermatitis and hockburn in broiler chicken flocks using optical flow, bodyweight and water consumption". In: *The Veterinary record of the British Veterinary Association* 180.20 (2017), p. 499.
- [10] Gerhard Manteuffel, Birger Puppe, and Peter C Schön. "Vocalization of farm animals as a measure of welfare". In: *Applied Animal Behaviour Science* 88.1 (2004), pp. 163–182. ISSN: 0168-1591. DOI: <https://doi.org/10.1016/j.applanim.2004.02.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0168159104000565>.
- [11] Sadia Naseem and Annie J King. "Ammonia production in poultry houses can affect health of humans, birds, and the environment-techniques for its reduction during poultry production". In: *Environmental science and pollution research international* 25.16 (June 2018), pp. 15269–15293. ISSN: 0944-1344. DOI: 10.1007/s11356-018-2018-y. URL: <https://doi.org/10.1007/s11356-018-2018-y>.
- [12] Suresh Neethirajan. "ChickTrack – A quantitative tracking tool for measuring chicken activity". In: *Measurement* 191 (2022), p. 110819. ISSN: 0263-2241. DOI: <https://doi.org/10.1016/j.measurement.2022.110819>. URL: <https://www.sciencedirect.com/science/article/pii/S0263224122001154>.
- [13] Raspberry Pi. *Raspberry Pi hardware*. <https://www.raspberrypi.com/>. Accessed: 7th May, 2024. 2024.
- [14] Lawrence Rabiner and Ronald Schafer. *Theory and Applications of Digital Speech Processing*. 1st. USA: Prentice Hall Press, 2010. ISBN: 0136034284.
- [15] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". In: *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*. Ed. by Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett. 2015, pp. 91–99. URL: <https://proceedings.neurips.cc/paper/2015/hash/14bfa6bb14875e45bba028a21ed38046-Abstract.html>.
- [16] Chris Stauffer and W. Eric L. Grimson. "Adaptive Background Mixture Models for Real-Time Tracking". In: *1999 Conference on Computer Vision and Pattern Recognition (CVPR '99), 23-25 June 1999, Ft. Collins, CO, USA*. IEEE Computer Society, 1999, pp. 2246–2252. DOI: 10.1109/CVPR.1999.784637. URL: <https://doi.org/10.1109/CVPR.1999.784637>.
- [17] Ultralytics. *ultralytics/YOLOv5*. <https://github.com/ultralytics/yolov5>. Accessed: 7th May, 2024. 2024.
- [18] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion". In: *J. Mach. Learn. Res.* 11 (2010), pp. 3371–3408. DOI: 10.5555/1756006.1953039. URL: <https://dl.acm.org/doi/10.5555/1756006.1953039>.
- [19] Xiao Yang, Lilong Chai, Ramesh Bahadur Bist, Sachin Subedi, and Zihao Wu. "A Deep Learning Model for Detecting Cage-Free Hens on the Litter Floor". In: *Animals* 12.15 (2022). ISSN: 2076-2615. DOI: 10.3390/ani12151983. URL: <https://www.mdpi.com/2076-2615/12/15/1983>.
- [20] Zoran Zivkovic. "Improved Adaptive Gaussian Mixture Model for Background Subtraction". In: *17th International Conference on Pattern Recognition, ICPR 2004, Cambridge, UK, August 23-26, 2004*. IEEE Computer Society, 2004, pp. 28–31. DOI: 10.1109/ICPR.2004.1333992. URL: <https://doi.org/10.1109/ICPR.2004.1333992>.

Topic Modeling of the SrpELTeC Corpus: A Comparison of NMF, LDA, and BERTopic

Teodora Mihajlov
0009-0008-8137-6750
Association for Language Resources
and Technologies
ul. Studentski trg 3, Belgrade, Serbia
Email: teodoramihajlov@gmail.com

Milica Ikonić Nešić
0000-0002-0835-8889
Univ. of Belgrade, F. of Philology
ul. Studentski trg 3, Belgrade, Serbia
Email: milica.ikonic.nesic@fil.bg.ac

Ranka Stanković, Olivera Kitanović
0000-0001-5123-6273
0000-0002-7571-2729
Univ. of Belgrade, F. of Mining and Geology
ul. Đušina 7, Belgrade, Serbia
Email: {ranka.stankovic, olivera.kitanovic}@rgf.bg.ac.rs,

Abstract—Topic modeling is an effective way to gain insight into large amounts of data. Some of the most widely used topic models are Latent Dirichlet allocation (LDA) and Nonnegative Matrix Factorization (NMF). However, new ways to mine topics have emerged with the rise of self-attention models and pre-trained language models. BERTopic represents the current state-of-the-art when it comes to modeling topics. In this paper, we compared LDA, NMF, and BERTopic performance on literary texts in the Serbian language, both quantitatively by measuring Topic Coherency (TC) and Topic Diversity (TD), and by conducting a qualitative evaluation of the obtained topics. Additionally, for BERTopic, we compared multilingual sentence transformer embeddings with the Jerteh-355 monolingual embeddings for Serbian. NMF yielded the best Topic Coherency results, while BERTopic with Jerteh-355 embeddings gave the best Topic Diversity. The monolingual Serbian Jerteh-355 embeddings also outperformed sentence transformer embeddings in both TC and TD.

Index Terms—topic modeling, LDA, NMF, BERTopic, SrpELTeC, computational literary studies

I. INTRODUCTION

TOPIC modeling has proven to be an effective tool for uncovering common themes and the underlying narratives in texts and for describing copious datasets. In social sciences, one way to leverage topic modeling is to explore topics in literary texts [1], [2].

In this paper, we present an evaluation of statistical and deep learning topic models on the SrpELTeC collection, Serbian part of ELTeC, the European Literary Text Collection, produced within COST Action CA16204 [3], [4]. The aims of the presented project are two-fold: (1) to explore topics in the SrpELTeC collection; (2) to evaluate and contrast the efficacy of conventional topic models, namely Latent Dirichlet Allocation (LDA) and Non-Negative Matrix Factorization (NMF), with a transformer-based topic model, BERTopic, in analyzing long texts in the Serbian language. To the best of our knowledge, BERTopic has only been used for modeling

topics in short texts in the Serbian language, namely on a dataset of tweets expressing hesitancy towards COVID-19 vaccination [5], where it outperformed both LDA and NMF. The model has not yet been applied to long or literary texts in the Serbian language.

To that end, this paper exploits natural language processing methods to obtain information about the Serbian literary texts, both in and outside of canon, which are presented in the SrpELTeC collection. The broader aim of this work is to, in the future, compile a comprehensive connected network of Serbian literary publications, based on the principles of Wikidata [6]. The insights that could further be derived from the project could be used not only for testing methods for modeling Serbian literary texts but also to pave the way for the development of approaches for digital humanities for Serbian. In the future, we will aim to rely on the principles of digital humanities that promote using big data paired with carefully curated metadata, following the example of the *MiMoText* research project in computational literary studies [7].

The remainder of the paper is structured as follows: section II delves into related work, covering both traditional and deep learning topic modeling methods and their use thus far; in III we cover data preprocessing steps (III-A) and methods used for obtaining text topics (III-B) and in IV we present the results, both qualitatively IV-A and qualitatively IV-B. Finally, in V, we lay out concluding remarks and propositions for further.

II. RELATED WORK

Several methods provide insight into latent topics in texts. Two of the most widespread methods for topic modeling are Latent Dirichlet Allocation (LDA) [8], and Nonnegative Matrix Factorization (NMF) [9]. LDA is a generative probabilistic model, specifically, a three-level Bayesian model, which models each item in a corpus as a representation of

probabilities of underlying topics [8]. It has proven to be one of the most popular topic modeling algorithms [10]. In contrast, NMF is a non-probabilistic linear, decompositional algorithm, which relies on matrix factorization [9]. In the context of topic modeling, NMF is based on TF-IDF, transforming data by breaking down a matrix into two lower-ranking matrices [10]. Both models call for a predetermined number of topics. Adjusting the number of topics and fitting other parameters accordingly can be challenging [11]. In addition, the models call for extensive data preprocessing. Another downside of traditional methods such as LDA and NMF is that they represent documents in a bag-of-words fashion, which ignores both word order and their semantic relationship [12].

In recent years, the rise of self-attention [13] paved the way for the development of pre-trained language models (PLMs). In turn, this facilitated generating word embeddings and adjusting them for different tasks, such as topic modeling. BERTopic is a BERT-based PLM trained on the topic modeling task, which utilizes pre-trained embeddings to generate text topics [14]. On top of the generated embedding, BERTopic leverages dimensionality reduction and clustering techniques, which are by default UMAP and HDBSCAN, respectively. To create topic representations, the model uses c-TF-IDF, a class-based variation of TF-IDF [10]. One of the perks of BERTopic is its modularity. Although the model has default settings for each of the aforementioned steps, the user can choose different algorithms and parameters for each of the steps, adjusting the model to their data and goals, which makes it a scalable topic modeling solution [14]. Unlike LDA and NMF, BERTopic does not require a predefined number of topics. The main downside of the model is that it assigns only one topic to each document [14].

III. MATERIALS AND METHODS

A. Data Description and Preprocessing

The data comprises the Serbian part of the ELTeC corpus - a multilingual collection of novels written in the period 1840-1920. The Serbian ELTeC collection, encompasses 100 novels, while the entire collection consists of 157 novels [3]. The 100 novels used here were written by 66 different authors, 62 male, and 4 female, and were published between 1852 and 1920. Two novels in the collection are written by unknown authors. The average novel length is 49,315 words. The remaining novels are currently being prepared and will be a part of an extended sub-collection SrpELTeC-ext.3 [3]. For the purpose of this research we used SrpELTeC TXM Copus¹ of 108 novels in level-2.² Novels in level-2 are annotated with part of speech (POS), lemma, and 7 categories of named entities: persons (PERS), organisations (ORG), locations (LOC), demonyms (DEMO), work of art (WORK), events (EVENT), and titles and professions (ROLE) [15]. Such annotated corpus allows for analysis to be conducted using only nouns (NOUN),

of which there are a total of 854,835 in this collection, with 30,684 being unique.

We used a spaCy Python package for the Serbian language for text preprocessing.³ First, we removed special characters and converted text to lowercase. Next, we converted the text from the Serbian Cyrillic script to Latin script. Finally, we tokenized and lemmatized the text, and removed stopwords. The stopwords consisted of a list of stopwords for the Serbian language, as well as corpus-specific stopwords. The corpus-specific stopwords were extracted by observing keywords while implementing the initial versions of all three models.

B. Models

Latent Dirichlet allocation (LDA). We implemented an LDA model using the Gensim library. Text tokens were obtained using the BoW approach, whereby both bigrams and trigrams were created. Subsequently, we generated a TF-IDF representation of the documents and filtered out words with a frequency < 0.03. To fine-tune the number of topics, we iterated the number of topics between 2 and 10 and evaluated Topic Coherence for each iteration. Finally, we picked 5 topics, since that yielded the best TC score.

Nonnegative Matrix Factorization (NMF). NMF was implemented with the Sci-kit Learn library, as it displayed significantly better results in comparison to its Gensim equivalent. The minimal word frequency was set to 15, and the maximal frequency was an occurrence of a word in 80% of the documents. A full list of stopwords, i.e. Serbian and corpus-specific, was passed to the model, in case some were not initially removed. After calculating topic coherence for the number of topics between 2 and 10, we set the parameter to 7 topics.

BERTopic. For BERTopic, we exploited the modular architecture of the model and tried to best fit each of its components to our data and research aims. We first generated word embeddings. As the default word embeddings, BERTopic uses sentence transformers [16], which support English and include three multilingual sentence transformer models that are trained for 50+ languages including Serbian:

- *distiluse-base-multilingual-cased-v2*: the model maps into 512-dimensional dense vector space and can be used for tasks like clustering or semantic search (480 MB, 135 million parameters).
- *paraphrase-multilingual-MiniLM-L12-v2*: this model maps sentences to a 384-dimensional dense vector space (420 MB, 117 million parameters).
- *paraphrase-multilingual-mpnet-base-v2*: this model maps sentences to a 768-dimensional dense vector space (970 MB, 278 million parameters).

In addition to the three multilingual embedding models, we tested the *Jerteh-355* embeddings. The *Jerteh-355* model is the largest model trained specifically for the Serbian language [17]. Although the model is not fine-tuned for the semantic search task, we wanted to see how it performs in

¹<https://live.european-language-grid.eu/catalogue/corpus/23621>

²<https://github.com/COST-ELTeC/ELTeC-srp/tree/master/level2>

³<https://github.com/procesaur/srpski>

TABLE I
TOPIC COHERENCE AND TOPIC DIVERSITY OF THE MODELS

Model	TC	TD
LDA	0.361	0.940
NMF	0.568	0.757
BERTopic		
<i>distiluse-base-multilingual-cased-v2</i>	0.427	0.869
<i>paraphrase-multilingual-MiniLM-L12-v2</i>	0.387	0.864
<i>paraphrase-multilingual-mpnet-base-v2</i>	0.299	0.925
<i>Jerteh-355</i>	0.456	0.970

TABLE II
LDA GENERATED KEYWORDS (NOUN)

Topic	Top 10 keywords
Topic0	knez (<i>knyaz</i>), vojvoda (<i>duke</i>), vojska (<i>army</i>), dvor (<i>castle</i>), gospodar (<i>lord</i>), junak (<i>hero</i>), car (<i>tsar</i>), šator (<i>tent</i>), pop (<i>priest</i>), vlastela (<i>Medieval Serbian nobility</i>)
Topic1	gospoda (<i>mam</i>), gospodin (<i>sir</i>), pop (<i>priest</i>), doktor (<i>doctor</i>), učitelj (<i>teacher</i>), škola (<i>school</i>), mati (<i>mother</i>), kapetan (<i>captain</i>), manastir (<i>monastery</i>), dete (<i>child</i>)
Topic2	narod (<i>people</i>), čovek (<i>man</i>), gospodar (<i>lord</i>), kapetan (<i>captain</i>), vezir (<i>vizier</i>), knez (<i>knyaz</i>), gospodin (<i>sir</i>), kmet (<i>serf</i>), čiča (<i>uncle</i>), koliba (<i>hut</i>)
Topic3	gazda (<i>lord</i>), gospodar (<i>sir</i>), pop (<i>priest</i>), čovjek (<i>man</i>), riječ (<i>word</i>), planina (<i>mountain</i>), dućan (<i>store</i>), talijer (<i>talir</i>), vrijeme (<i>time/weather</i>), djeca (<i>children</i>)
Topic4	društvo (<i>society</i>), čovek (<i>man</i>), reč (<i>word</i>), deca (<i>children</i>), načelo (<i>principle</i>), sloboda (<i>freedom</i>), dete (<i>child</i>), stanje (<i>condition</i>), ženskinja (<i>woman</i>), nauka (<i>science</i>)

comparison to the aforementioned multilingual models. The model specifics are as follows:

- *Jerteh-355* the model size is 355 million parameters, and it was trained on 4 billion tokens in the Serbian language.

For the dimensionality reduction step, we used UMAP, with the following parameters - $n_neighbors = 5$; $n_components = 5$. For clustering, we used HBDSCAN with the minimal cluster size set to 3. The rest of the UMAP and HBDSCAN parameters were default.

Lastly, to create topic representations, BERTopic utilizes CountVectorizer and class-based c-TF-IDF, to model the importance of each document cluster. We used CountVectorizer to filter out noise from the data: additional stopwords were removed, and all words with frequency <5 and $>80\%$ of the documents were filtered out. We looked at both unigrams and bigrams.

After generating topics, BERTopic creates a -1 topic that contains outlier documents. To remove outliers, it offers an *outlier_reduction* function. However, when we tried using this function, we got different topic keywords, which were worse than those originally generated. Therefore, in this phase of research, we opted out of using the *outlier_reduction* option.

C. Evaluation

The models were evaluated quantitatively and qualitatively. For the quantitative evaluation, we used Topic Coherence

(TC) and Topic Diversity (TD) measures, both of which are frequently used for evaluating topic models [14], [10], [18]. TC is a measure of semantic relatedness between the words for each topic [19]. We used the C_V coherence measure, which is based on a combination of a sliding window, a one-set segmentation of the top words, and an indirect confirmation measure that uses normalized pointwise mutual information (NPMI) and the cosine similarity [20]. TC ranges from 0 to 1, with values closer to 1 signify more related topic words. TD computes the percentage of unique words for top_n words for each topic. It ranges from 0 to 1, where 1 marks more related, and 0 more redundant topics [18].

IV. RESULTS AND DISCUSSION

A. Quantitative Evaluation

Using all POS, NMF significantly outperformed both LDA and BERTopic in Topic Coherency (TC). BERTopic, however, generated the most diverse topics (Topic Diversity, TD). In topic diversity, LDA came close to BERTopic, while NMF displayed a significant difference in TD measure in comparison with the two other models, as displayed in Table I. The best model performance is presented in bold, while the best performance among different embeddings for BERTopic is presented in bold and underlined font.

For BERTopic, we can see that, although it was not fine-tuned for semantic search, a monolingual embedding model for Serbian, *Jerteh-355* outperformed the three sentence transformer models in both TC and TD (Table I).

B. Qualitative Evaluation

For the qualitative analysis of the obtained topics, we will look into *top_n* keywords for each of the models. Translations of all keywords are presented in brackets. Personal names

TABLE III
NMF GENERATED KEYWORDS (NOUN)

Topic	Top 10 keywords
Topic0	vojska (<i>army</i>), drum (<i>road</i>), vojnik (<i>soldier</i>), kapetan (<i>captain</i>), borba (<i>battle</i>), neprijatelj (<i>enemy</i>), bol (<i>pain</i>), komanda (<i>comand</i>), oficir (<i>officer</i>), planina (<i>mountain</i>)
Topic1	knez (<i>knyaz</i>), đeneral (<i>general</i>), gospodar (<i>sir</i>), ministar (<i>minister</i>), seljak (<i>peasant</i>), načelnik (<i>chief</i>), otrov (<i>poison</i>), doktor (<i>doctor</i>), vojvoda (<i>duke</i>), svetlost (<i>duke</i>)
Topic2	despot (<i>despot</i>), vojska (<i>army</i>), vojvoda (<i>duke</i>), grad (<i>city</i>), paš (<i>pasha</i>), sultan (<i>Sultan</i>), kaluder (<i>monk</i>), car (<i>emperor</i>), dvor (<i>castle</i>), manastir (<i>monastery</i>)
Topic3	slovo (<i>letter</i>), đak (<i>pupil</i>), učitelj (<i>teacher</i>), manastir (<i>monastery</i>), kmet (<i>serf</i>), iguman (<i>Hegumen</i>), majstor (<i>meister</i>), gazda (<i>lord</i>), kaluder (<i>monk</i>), arhimandrit (<i>archimandrite</i>)
Topic4	vezir (<i>vizier</i>), ratnik (<i>warrior</i>), hajduk (<i>hajduk</i>), aga (<i>agha</i>), vojvoda (<i>duke</i>), družina (<i>length</i>), gospodar (<i>lord</i>), tvrđava (<i>fortress</i>), straža (<i>watch, as in Night watch</i>), grad (<i>city</i>)
Topic5	gospoda (<i>mam</i>), doktor (<i>doctor</i>), kćer (<i>daughter</i>), gospodar (<i>lord</i>), udovica (<i>widow</i>), gospođica (<i>mam</i>), sahat (<i>hour</i>), tetak (<i>uncle</i>), advokat (<i>lawyer</i>), dama (<i>lady</i>)
Topic6	pop (<i>priest</i>), kmet (<i>serf</i>), učitelj (<i>teacher</i>), kapetan (<i>captain</i>), baba (<i>grandmother</i>), gospoja (<i>lady</i>), sokak (<i>street</i>), avlija (<i>courtyard</i>), čata (<i>clerk</i>), gazda (<i>lord</i>)

TABLE IV
BERTOPIC GENERATED KEYWORDS BY THE *jerteh-355* MODEL EMBEDDINGS (NOUN)

Topic	No. of Documents	Top keywords
-1	9	despot (<i>despot</i>), kir (<i>lord</i>), sultan (<i>sultan</i>), česar (<i>emperor</i>), vezir (<i>vizier</i>), bula (<i>Muslim woman</i>), knežević (<i>prince</i>), vlastelin (<i>nobleman</i>), patrijarh (<i>patriarch</i>), vjera (<i>faith</i>)
0	19	hanum (<i>lady</i>), fratar (<i>friar</i>), robinja (<i>slave woman</i>), gospoja (<i>madam</i>), hanuma (<i>lady</i>), tatko (<i>father</i>), naprednjak (<i>progressive</i>), đakon (<i>deacon</i>), fala (<i>thanks</i>), nena (<i>grandmother</i>)
1	10	arhimandrit (<i>archimandrite</i>), major (<i>major</i>), patrijarh (<i>patriarch</i>), grof (<i>count</i>), nastojatelj (<i>superior</i>), djevojka (<i>girl</i>), koi (<i>which</i>), zadruga (<i>cooperative</i>), cigareta (<i>cigarette</i>), riječ (<i>word</i>)
2	7	đakon (<i>deacon</i>), monah (<i>monk</i>), kasta (<i>caste</i>), frajla (<i>lady</i>), ogrlica (<i>necklace</i>), predsednik (<i>president</i>), čika (<i>uncle</i>), forinta (<i>forint</i>), senat (<i>senate</i>), adidar (<i>jewelry</i>)
3	6	aga (<i>aga</i>), subaša (<i>overseer</i>), čorbadžić (<i>chief</i>), kahva (<i>coffee</i>), hodža (<i>imam</i>), tatko (<i>father</i>), loža (<i>lodge</i>), duhovnik (<i>spiritual father</i>), riječ (<i>word</i>), svijet (<i>world</i>)
4	6	nazaren (<i>nazarene</i>), bukvar (<i>primer</i>), gradina (<i>garden</i>), tablica (<i>tablet</i>), tabla (<i>board</i>), cigla (<i>brick</i>), slovo (<i>letter</i>), apostol (<i>apostle</i>), sotona (<i>Satan</i>), crep (<i>roof tile</i>)
5	6	ćata (<i>boss</i>), aga (<i>aga</i>), gospa (<i>lady</i>), front (<i>front</i>), stanica (<i>station</i>), baterija (<i>battery</i>), vagon (<i>wagon</i>), dućandžija (<i>shopkeeper</i>), divizija (<i>division</i>), automobil (<i>car</i>)
6	5	šator (<i>tent</i>), arhimandrit (<i>archimandrite</i>), trpezarija (<i>dining room</i>), vojvoda (<i>duke</i>), knez (<i>prince</i>), vlastela (<i>nobility</i>), tag (<i>tag</i>), prisednik (<i>president</i>), beležnik (<i>notary</i>), društvo (<i>society</i>), knez vojvoda (<i>duke prince</i>)
7	5	grofica (<i>countess</i>), kneginjica (<i>princess</i>), dragana (<i>darling</i>), grof (<i>count</i>), nana (<i>grandmother</i>), urednik (<i>editor</i>), teta (<i>aunt</i>), drama (<i>drama</i>), gospa (<i>lady</i>), šor (<i>street</i>)
8	5	vezir (<i>vizier</i>), gospodin ministar (<i>minister</i>), vranac (<i>black horse</i>), kavana (<i>tavern</i>), žandarm (<i>gendarme</i>), posluživanje (<i>service</i>), česma (<i>fountain</i>), aga (<i>aga</i>), pobra (<i>peasant</i>), glavlar (<i>chief</i>)

are not translated but are capitalized and indicated with an abbreviation *pers.*. Labels of keywords relating to other named entities such as cities or buildings are also presented in *italic*, preceding word translation.

Analysis of the keywords of the final LDA model presented in Table II determined that the model generated diverse keywords.

However, many keywords are names of the characters in the novels. Even though they are relevant to the corpus, they do not tell us much about latent topics in the texts, especially if one is unfamiliar with novels included in SrpELTeC. By carefully examining NMF keywords (Table III), we ascertained that the model yielded the most informative keywords, which refer to the topics discussed in the novels. The keywords mention words such as **turčin** (eng. Turk), **vojvoda** (duke), **manastir** (eng. monastery),

combined with character names. The keywords are informative of both main topics discussed in the novels and of the main characters that represent the corpus.

No matter the embedding model or parameter setting, BERTopic continually generated names of the characters in the novels as keywords (Table IV), which do not provide us with sufficient information about the topics. After adjusting the model parameters, only 6 documents out of 100 (6%) were assigned the -1 topic, i.e. classified as outliers.

V. CONCLUSION AND FURTHER WORK

In this paper, we compared traditional topic models, NMF and LDA, with a transformer-based BERTopic, on the SrpELTeC collection. Although we expected BERTopic to outperform the traditional models significantly, it only did so when it came to topic diversity. When it comes to TC, it fell somewhat short compared to NMF.

To address the limitations of current work, we plan to try and see the effects of using chunked text on topic extraction.

Even though our current corpus comprises 49,315 words, it consists of just 100 documents. Both the length of the texts as well as the small number of documents could be hindering the current model's performance. Therefore, better topics might be obtained by passing shorter chunks to BERTopic might improve results. Lastly, we plan on trying to improve BERTopic performance with chunked documents, as well as see how it performs with other embedding models for the Serbian language, such as BERTić [21].

ACKNOWLEDGMENT


This research was supported by the Science Fund of the Republic of Serbia, #7276, *Text Embeddings - Serbian Language Applications - TESLA*.

REFERENCES

- [1] I. Uglanova and E. Gius, "The order of things. a study on topic modelling of literary texts." *CHR*, no. 18-20, p. 2020, 2020.
- [2] K. E. Chu, P. Keikhosrokiani, and M. P. Asl, "A topic modeling and sentiment analysis model for detection and visualization of themes in literary texts," *Pertanika Journal of Science & Technology*, vol. 30, no. 4, pp. 2535–2561, 2022, <https://doi.org/10.47836/pjst.30.4.14>.
- [3] R. Stanković, C. Krstev, B. Š. Todorović, D. Vitas, M. Škorić, and M. I. Nešić, "Distant reading in digital humanities: Case study on the serbian part of the eltec collection," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, 2022, pp. 3337–3345. [Online]. Available: <https://aclanthology.org/2022.lrec-1.356>
- [4] C. Schöch, T. Erjavec, R. Patras, and D. Santos, "Creating the european literary text collection (eltec): Challenges and perspectives," *Modern Languages Open*, 2021, <http://doi.org/10.3828/mlo.v0i0.364>.
- [5] D. Medvečki, B. Bašaragin, A. Ljajić, and N. Milošević, "Multi-lingual transformer and bertopic for short text topic modeling: The case of serbian," in *Conference on Information Technology and its Applications*. Springer, 2024, pp. 161–173, https://doi.org/10.1007/978-3-031-50755-7_16.
- [6] D. Vrandečić and M. Krötzsch, "Wikidata: a free collaborative knowledgebase," *Communications of the ACM*, vol. 57, no. 10, pp. 78–85, 2014, <https://doi.org/10.1145/2629489>.
- [7] C. Schöch, M. Hinzmann, J. Röttgermann, K. Dietz, and A. Klee, "Smart modelling for literary history," *International Journal of Humanities and Arts Computing*, vol. 16, no. 1, pp. 78–93, 2022, <https://doi.org/10.3366/ijhac.2022.0278>.

- [8] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [9] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," *Advances in neural information processing systems*, vol. 13, 2000. [Online]. Available: <https://api.semanticscholar.org/CorpusID:2095855>
- [10] R. Egger and J. Yu, "A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts," *Frontiers in sociology*, vol. 7, p. 886498, 2022.
- [11] —, "Identifying hidden semantic structures in instagram data: a topic modelling comparison," *Tourism Review*, vol. 77, no. 4, pp. 1234–1246, 2021.
- [12] M. Švaňa, "Social media, topic modeling and sentiment analysis in municipal decision support," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 2023, pp. 1235–1239, <http://dx.doi.org/10.15439/2023F1479>.
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017, <https://doi.org/10.48550/arXiv.1706.03762>.
- [14] M. Grootendorst, "Bertopic: Neural topic modeling with a class-based tf-idf procedure," *arXiv preprint arXiv:2203.05794*, 2022, <https://doi.org/10.48550/arXiv.2203.05794>.
- [15] R. Stanković, C. Krstev, B. Šandrih Todorović, and M. Škorić, "Annotation of the serbian eltec collection," *Infotecha - Journal for Digital Humanities*, vol. 21, no. 2, pp. 43–59, 2022. [Online]. Available: https://infoteka.bg.ac.rs/ojs/index.php/Infoteka/article/view/2021.21.2.3_en
- [16] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*, 2019, <https://doi.org/10.48550/arXiv.1908.10084>.
- [17] M. Škorić, "Novi jezički modeli za srpski jezik," *Infoteka*, vol. 24, 2024, <https://doi.org/10.48550/arXiv.2402.14379>. [Online]. Available: <https://arxiv.org/abs/2402.14379>
- [18] A. B. Dieng, F. J. Ruiz, and D. M. Blei, "Topic modeling in embedding spaces," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 439–453, 2020, <https://doi.org/10.48550/arXiv.1907.04907>.
- [19] D. Newman, J. H. Lau, K. Grieser, and T. Baldwin, "Automatic evaluation of topic coherence," in *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, 2010, pp. 100–108.
- [20] M. Röder, A. Both, and A. Hinneburg, "Exploring the space of topic coherence measures," in *Proceedings of the eighth ACM international conference on Web search and data mining*, 2015, pp. 399–408, <https://doi.org/10.1145/2684822.2685324>.
- [21] N. Ljubešić and D. Lauc, "BERTić - the transformer language model for Bosnian, Croatian, Montenegrin and Serbian," in *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*. Kiyv, Ukraine: Association for Computational Linguistics, Apr. 2021, pp. 37–42, <https://doi.org/10.48550/arXiv.2104.09243>. [Online]. Available: <https://www.aclweb.org/anthology/2021.bsnlp-1.5>

Towards to an interface design for a building operations CPS

Filipe Moreira^{1*}, Rosana Alexandre¹, Rosa Silva¹, Jo^ao Oliveira¹, Manuel Alves¹,
Jo^ao Pereira², Ana Colim¹, Nelson Rodrigues¹ 

¹ DTx – Digital Transformation CoLAB, University of Minho, Portugal, e-mail {filipe.moreira, ana.colim}@dtx-colab.pt

² dst group – Rua de Pitancinhos - Palmeira 4700-727 Braga, Portugal, e-mail joaodaniel@innovpoint.com

Abstract — Cyber-physical systems (CPS) enable value creation in software development. However, they present several challenges that necessitate a structured approach to design and implementation. We present a concept on how to approach the design and development of a software service, based on multiple CPS. It delineates the specific characteristics of a CPS, including graphical user interface (GUI) components and architectural aspects, while recommending appropriate tools and design methodologies. This approach aligns with contemporary frameworks such as Industry 4.0 and Software as a Service (SaaS). The multidisciplinary nature of CPS leads to several development challenges, and we focus our work on using design thinking techniques, methodologies, and tools that are valuable solutions and key to the sustainability of new designs. These insights may help researchers and industrial practitioners to develop and commercialize service-based CPS.

Index Terms — Cyber-Physical Systems; Industry 4.0; User Support; Graphical User Interface, Design Thinking, Building Operation.

I. INTRODUCTION

The development and integration of industrial Cyber-Physical Systems (CPS), systems that integrate computation and physical processes [8], is a topic of utmost importance, given the significant challenges it presents. The severe consequences of system failures, which can impact safety, productivity, costs, and company reputation, underscore the critical nature of this field. Traditionally, components of CPS, such as physical assets, digital twins, industrial applications, and critical technological aspects like communication and cybersecurity, are developed independently. This traditional siloed approach poses a significant hurdle in harmonizing diverse development paradigms to deliver a seamless user experience while minimizing research and development costs.

Digital twins, unique and essential elements within CPS, are defined variably across literature but represent a digital replica of a physical asset [6]. They support various industrial applications through accessible data and services, operating on local servers or via cloud computing. The emphasis on user-centric design enhances value to industrial practitioners. The integration of digital twins requires a structured approach to mirror the connectivity of physical entities and ensure interoperability across different vendors and building operations stakeholders. The paper outlines the complexities

of designing CPS by discussing integrating systems and subsystems into a unified framework that supports robust data management and system operation across different building operation phases, under the design perspective. This technical integration extends to aligning business models and ensuring system components deliver continuous value, adhering to traditional and emerging business practices. Moreover, this paper underscores the integration of human factors, recognizing the importance of seamless interaction between users and CPS. It involves developing adaptable interfaces and fostering safe and efficient human-machine collaboration, an essential aspect of CPS.

II. METHODOLOGY

This study is part of a project focused on developed by our design team within a project focused on the development of graphical user interfaces for software-as-a-service (SaaS) within a smart-building management operations context.

The design team's main challenge was articulating the functional structure, of the several CPS, in a consistent, intuitive and pleasant user experience.

The project involves is a well-established and accomplished construction company in Northern Portugal with an extensive client portfolio. This project is part of a significant modular construction program that benefits from all the newest digital technologies a CPS can offer. Our approach consisted of using design thinking and systems thinking methodologies together to better address the complexity-related problems during CPS design and implementation [4].

To effectively execute this project, a multidisciplinary team established for this project involves distinct groups: industrial practitioners, a team of researchers from DTx's Ergonomics and Engineering Psychology team, which includes two designers, the Edge and Embedded Computing group, and the DTx Data and Application Engineering group. Moreover, the structure of graphical interfaces should be agnostic enough to adapt to these diverse types of buildings, data collected, and

* Corresponding Author's email: filipe.moreira@dtx-colab.pt

operations performed by users with different profiles and access permissions. The existence of a document with wireframe characteristics of the software provided by the project's leading company allowed the requirements-gathering work to progress more quickly. The leading company in the project will play a dual role, acting both as a vendor and as a client of the final SaaS that will be born from the developed prototypes. This unique position underscores the company's responsibilities and the need for a comprehensive understanding of the system's design and functionality.

From the beginning, it was clear that the terminology used required a phase of learning and studying the context. This required meetings to discuss and delve into the concepts reflected in this document. Clarifying any misunderstandings needed to be considered. These meetings were crucial in preventing the implementation of interfaces and back-end software that might need restructuring later.

The meetings also fostered a collaborative environment, allowing an understanding of the document's author, including his motivations, professional context, and business ideas. This insight improved the ability to reference a set of conditions, enhanced communication, and contributed to delivering a user-focused project.

After the first meetings to address the issues raised by the document with the project leader's wireframe, our work consisted of prioritizing some activities, which we will explore in more detail in this article. These activities included:

- a) Creating a project glossary, where we could find a common language free from misunderstandings.
- b) Mapping user stories of all potential users of the system, listing the accessible functionalities and activities of each one within the system.
- c) Developing a low-fidelity prototype, where we organized the functional hierarchies into a structure of standard interfaces.
- d) Crafting a service design blueprint, to uncover touchpoints that the user has with the service.
- e) Articulating different users involved in the building maintenance ticketing process.
- f) Creating a medium-fidelity prototype, with the fine-tuning of details and interaction mechanics per user, anticipating and paving the way for high-fidelity implementation and prototyping in Front-End.

Figure 1 illustrates the development stages of the main phases of the framework to support the design process. As mentioned previously, these five activities were carried out by the design team. Despite appearing separately, the process often occurred simultaneously and iteratively, as layers of information were added.

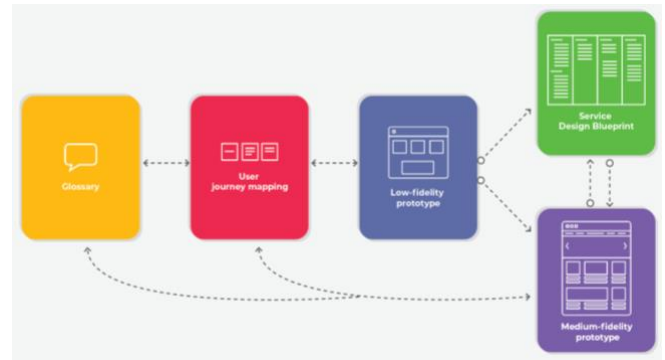


Fig 1: Representation of the main phases to support the design process

III. RESULTS

From a human-centered design perspective, the design team conducted co-creation sessions with the stakeholders and the development team to understand users' needs and potential solutions [10].

A. Glossary

Given the complexities of the software and hardware platforms that make up this CPS, the initial focus of the interface design project was to initiate the development of a glossary that would allow discussing different views of the nomenclatures used, allowing for unification in the language used by the team multidisciplinary involved.

The creation of the glossary began with an in-depth reading of the initial wireframe, looking for inconsistencies in the terms used, discussing their meanings with the different stakeholders of the teams involved in the project, suggesting new nomenclatures, listing the different definitions and applications alphabetically, and considering their context and usage.

These discussions with the teams were carried out throughout meetings to approach the content of this document, simultaneously with the discussion of the organizational structure of the developing software and the identification of each of the potential users of the platform. In an effort to achieve gradual and sustained discovery, the goal was to identify a practical and optimal logic that will enable straightforward and clear use of its functionalities.

The determination of some of the nomenclatures, evident at the beginning for users with an elevated level of expertise, were gradually adapted whenever we found potential conflicts between interpretations, considering the different degrees of specialization and knowledge of the users who will use this SaaS. The effort allocated to the articulation between the different teams made it possible to gather consensus and compromise solutions to better create well-structured software adaptable to different user profiles.

This search for a common language, applicable both to the scope of the teams involved in the development of this CPS

and to the scope of the different users who would operate as a team in its use, was fundamental to emphasize the social and human dynamics that should be an integral part of a CPS under construction [13].

B. User Story Maps

Another tool for the interface design was the application of User Story Maps. This technique allows for identifying what is being developed from the perspective of the end users and helps to understand the reasoning behind the solution. It provides an overview of the product/service while also allowing it to be visualized in a fragmented manner [11].

When building the user story map, the team identified the different contexts of use for each potential user of the CPS, the functionalities that each can and should access, and the access permissions at various levels of specification of each component. More specifically, this resulted in a visual representation of the users of the CPS, with a hierarchy of core activities to be performed in the system and subsequent tasks.

The original wireframe that served as the starting point for this exercise suffered from the bias of being built from the perspective of a user with Super Administrator privileges, which is why the need to analyze the role played by each of the potential users of the software became critical. **This analysis enabled** the team to structure more holistically the different types of interfaces necessary and applicable to each user profile.

At the top of the hierarchical structure of the system is the Super Administrator, who has access to all its functionalities: capable of editing and parameterizing the system, creating

customer and user profiles stipulating roles and permissions, as well as creating, editing, and parameterizing all assets and equipment included in the system (Figure 2). All other users will have lower access levels to the system, depending on their role and the type of permissions assigned (Management Areas Interfaces, Maintenance Tasks Interface, or Final User/Client Interface).

The initial view of the software interface architecture, from the perspective of a Super Administrator, brought together all the accessible information of the recommended CPS, categorizing specificities and functionalities in a highly partial usage logic. Given the predictability small number of Super Administrator profiles, this view had limitation despite the breadth of processed, processable, and specifiable data.

Thus, a vision of a service-oriented architecture was required, making it more suitable for subsequent incorporations of new services and devices, commonly found in contexts of Smart Home environments. Such environments allow automation, emergency detection, and energy savings, among other features [12].

This Cyber-Physical Systems will have the functions of collecting and processing data to support the creation or maintenance of the building throughout its users. The structuring of data exchanges within the network of connected intelligent artifacts allows users to become its agents. Therefore, it is crucial to understand the central role of human beings in the complex organizational context of this case study. **Consequently**, the development of interfaces must respect the premise that the cyber-physical space must be designed and programmed by the social system presented in a work organization model.

Moreover, the human-machine relationship should create a potential symbiosis between the physical space (machines,

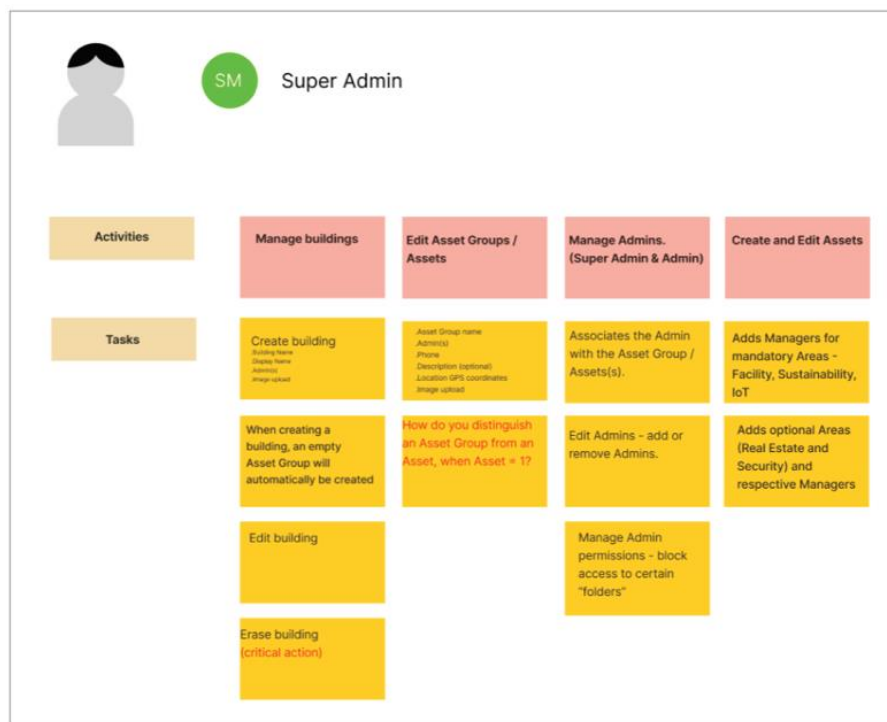


Fig. 2: Super Admin User story map

tools) and the social space (identities, mental references, decision-making, communication language, and organizational roles) through the effective design of its interfaces, with a clear objective of increasing productivity. Good interface design should allow increasing possibilities for intuitive interaction between humans and machines, facilitating operators' working conditions and improving the user experience in complex environments [3].

To this extent, using User Story Mapping makes the workflow or value chain visible, highlighting the relationships between different priority-level tasks, helping to confirm the integrity of the backlog, and providing a helpful context for prioritizing features for implementation. "Mapping your story helps you find holes in your thinking" [11, p. 11].

Additionally, user story maps facilitate productive, human-centered discussions about product creation, where it is possible to establish priorities regarding its activities and tasks. They provide a holistic view of the product, revealing logical and releasable slices of product increments that align with users' needs. Importantly, they also uncover potential impacts and areas of risk, allowing for early learning and decision-making. This early learning is a valuable tool for teams, enabling them to understand what works and what does not. Armed with this knowledge, teams can make informed decisions about where to focus first, maximizing usability, value, and feasibility in subsequent iterations. Furthermore, user story mapping can be used both as a form of documentation and communication between teams [9].

C. Low-fidelity Prototype and Service Design Blueprint

The following exercise was to create a low-fi prototype in Figma (<https://www.figma.com/>) where we sketched a set of

screens used in the initial and exploratory phases of the project and a tool that helps in requirements gathering [2]. We established the various stages of interaction between each user and the software, starting from a login screen and trying to understand the natural sequences of the availability of different functionalities. During this exercise, since the IoT Hub (where all the data collected by the different equipment and sensors will flow) was still in the study phase, it was understood that it was essential to focus on the Super-Administrator's perspective.

The Super-Administrator, as the key figure in our system, articulates all the platform's capabilities and makes them accessible on the interfaces of this privileged user. The idea that proved to be most interesting advocated that all other users' interfaces should correspond to the same graphical structure, sharing common interaction elements, which would adopt more simplified profiles depending on the limited access permissions of each user.

Throughout the process, given the questions raised during the meetings to present the low-fi prototype to the project leader, numerous technical issues were triggered, leading to revisions of the original wireframe.

We understood at this stage that it would be beneficial, within the scope of analyzing the flow of information throughout the process of generating tickets for building maintenance operations, to use a tool such as Service Design Blueprint (Figure 3), allowing a reading of the flow of procedures of each actor in the ticketing process, organizing them along a timeline and identifying the points of contact between them and potential points of friction or gaps in development and planning at this stage.

Given the intangible nature of services and their complexity, discussing them can be challenging. To this



Fig. 3: Service design blueprint for the ticketing process

extent, blueprinting helps to create a visual representation of the service process that highlights the steps in the process, the touchpoints that occur and the physical evidence that exists from the point of view of the customer and the employees involved [5], [7]. Blueprinting helps an organization identify points of failure, improvement areas, innovation opportunities, and opportunities to increase profit. It allows for a standard visualization tool to identify and demonstrate how a service currently works or how to design a new service process. “Service blueprinting allows the capturing of dynamic processes in a visual manner. Relatively few methods allow for this type of dynamic, and at the same time visual, representation” [1, p. 92].

This tool demonstrated its usefulness by highlighting the need for structural changes to the interface design, introducing new users with their specificities regarding the software’s use. Identifying these changes in this initial phase of the project made it possible to prevent development and implementation work that would most likely prove redundant or even useless in later phases of the project.

These changes were introduced in the different iterations of the execution of this prototype, fueling the discussion and providing new insights into the expected functioning of the software, motivating a continuous search for solutions reflected in the tasks of back-end developers.

D. Mid-Fidelity Prototype

When we considered the structure defined in the low-fidelity prototype satisfactorily completed, articulating the different actors with their interfaces throughout the numerous procedures that the software would allow, including ticket management tasks and maintenance teams, creation of new user profiles, new assets, and identification of smart devices, we fine-tuned the interaction mechanics, namely: buttons, breadcrumbs, headers, footers, side menus, iconography, color schemes, font size, spacing grid and hierarchical articulation of the various graphic components.

The adjustments at this stage allow for a more in-depth assessment of what could become the final product from the end user’s perspective. Discussing this Figma prototype with the project leader, the questions that arise essentially concern preferences, personal tastes, suggestions for minor changes, identification of “nice-to-have” features, the refinement of aesthetic characteristics, and, as a rule, reduced impact on the final product’s structure.

The maturity level of the medium-fidelity prototype advances with each iteration, something that we were able to observe throughout the presentation meetings of the work carried out with the project leader, paving the way in a safe and sustained way for the implementation phase that will follow, translating this prototype into code, enabling us to better assess its integration with the Backend team’s work and the implementation of the necessary infrastructures for the final prototype’s proper functioning.

This project aims to create a prototype very close to a final product, making it sufficiently solid and configurable so that the necessary adaptations for the project leader’s placement on the market are aligned with the current state-of-the-art web applications that are easily accessible by any type of common digital device (tablet, smartphone, or laptop).

IV. CONCLUSION

The pervasive existence of CPSs in today’s everyday world is an unavoidable reality. Human beings’ dependence on complex systems that facilitate access to information from the digital and physical world and subsequent decision-making has grown in line with the certainties and security that rapid technological evolution has allowed.

It is essential to maintain sight of the design purpose of these kinds of systems, created by man with the tools he finds at his disposal, to use them as new instruments of decision, action, and reaction. Given its importance as the ultimate recipient of use, the human variable in cyber-physical systems demands to be placed first on the list of priorities. Applying human-centred design methodologies, using design thinking techniques and best practices, allows a deep understanding of the users and their motivations, preferences, constraints, and potential for involvement as an active element in the system.

The human variable must be meticulously considered whenever a cyber-physical system interfaces with humans. Given its natural tendency to increase entropy, the optimal design of interfaces becomes of utmost importance in any CPS project. This aspect cannot be overlooked, highlighting the urgency of its consideration in system development.

ACKNOWLEDGMENT

This article is a result of the Innovation Pact “R2UTechnologies | modular systems” (C644876810-00000019), by “R2UTechnologies” Consortium, co-financed by NextGeneration EU, through the Incentive System “Agendas para a Inovação Empresarial” (“Agendas for Business Innovation”), within the Recovery and Resilience Plan (PRR).

REFERENCES

- [1] Bitner, M., Ostrom, A., & Morgan, F. (2008) California Management Review Service Blueprinting: A Practical Technique for Service Innovation, p.92. <https://doi.org/10.2307/41166446>
- [2] Dam, R. F. and Teo, Y. S. (2020, July 1). What Kind of Prototype Should You Create? Interaction Design Foundation - IxDF. <https://www.interaction-design.org/literature/article/what-kind-of-prototype-should-you-create>
- [3] El-Haouzi, H. B., Valette, E., Krings, B. J., & Moniz, A. B. (2021). Social dimensions in cps & iot based automated production systems. *Societies*, 11(3). <https://doi.org/10.3390/soc11030098>
- [4] Gürdür, D., & Törngren, M. (2018). Visual Analytics for Cyber-physical Systems Development: Blending Design Thinking and Systems Thinking. Presented at the 15th Annual NordDesign Conference (NordDesign 2018). Retrieved from <https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-232358>

- [5] Interaction Design Foundation - IxDF. (2016, June 3). What are Service Blueprints?. Interaction Design Foundation - IxDF. <https://www.interaction-design.org/literature/topics/service-blueprint>
- [6] Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the Digital Twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36–52. <https://doi.org/10.1016/j.cirpj.2020.02.002>
- [7] Kalbach, J. (2016). *Mapping experiences*. O'Reilly Media.
- [8] Lee, E. A. (2008, May). Cyber physical systems: Design challenges. In 2008 11th IEEE international symposium on object and component-oriented real-time distributed computing (ISORC) (pp. 363-369). IEEE.
- [9] Milicic, A., El Kadiri, S., Perdikakis, A., Ivanov, P., & Kiritsis, D. (2014). Toward the definition of domain concepts and knowledge through the application of the user story mapping method. *International Journal of Product Lifecycle Management*, 7(1), 3-16.
- [10] Norman, D.A.: *The Design of Everyday Things (Revised an)*. Basic Books (2013).
- [11] Patton, J., & Economy, P. (2014). *User story mapping: discover the whole story, build the right product*. " O'Reilly Media, Inc."
- [12] Romero, D., Hermosillo, G., Taherkordi, A., Nzekwa, R., Rouvoy, R., & Eliassen, F. (2013). The DigiHome service-oriented platform. *Software - Practice and Experience*, 43(10), 1205–1218. <https://doi.org/10.1002/spe.1125>
- [13] Wang, F.-Y. (2010). *The Emergence of Intelligent Enterprises: From CPS to CPSS*. www.computer.org/intelligent.

An environment model in multi-agent reinforcement learning with decentralized training

Rafał Niedziółka-Domański

0009-0004-0405-5508

Maria Curie-Skłodowska University
in Lublin

5 M. Curie-Skłodowskiej Square, 20-031 Lublin, Poland

Email: rniedziolkad@gmail.com

Jarosław Bylina

0000-0002-0319-2525

Maria Curie-Skłodowska University
in Lublin

5 M. Curie-Skłodowskiej Square, 20-031 Lublin, Poland

Email: jaroslaw.bylina@umcs.lublin.pl

Abstract—In multi-agent reinforcement learning scenarios, independent learning, where agents learn independently based on their observations, is often preferred for its scalability and simplicity compared to centralized training. However, it faces significant challenges due to the non-stationary nature of the environment from each agent’s perspective.

We investigate if incorporating an environment model in multi-agent reinforcement learning with decentralized training can alleviate the non-stationarity effects caused by the adaptive behaviors of other agents. To do this, we design and implement a custom model-based algorithm and compare its performance with the well-known model-free algorithm (Deep Q-Network). Our algorithm uses an environment model to plan and select actions. However, we do not require the model to be perfect for action selection, allowing it to be learned and improved during training. Our results suggest that integrating environment models into MARL offers a viable solution to mitigate non-stationarity.

Keywords: reinforcement learning, multi-agent reinforcement learning, model-based RL.

I. INTRODUCTION

IN MULTI-AGENT reinforcement learning (MARL), agents learn to make decisions in an environment where other agents are also learning and adapting. A major challenge in MARL is non-stationarity, where the perceived environment’s dynamics change from the perspective of each agent due to the adaptive behaviors of other agents. This non-stationarity can significantly hinder the learning process, making it difficult for agents to converge on optimal strategies.

Decentralized learning, where agents learn independently based on their observations, is often preferred for its scalability and simplicity. However, it faces significant challenges due to the non-stationary nature of the environment from each agent’s perspective. This article explores whether using an environment model in MARL with independent learning can mitigate the non-stationarity problem caused by other adaptive agents.

To investigate this, we design and implement a custom model-based method for MARL. We then compare the performance of our model-based algorithm against a traditional model-free algorithm to assess its effectiveness in reducing non-stationarity.

Our experiments and results provide insights into the potential benefits of integrating environment models in decentralized

MARL, offering an alternative perspective on addressing the non-stationarity challenge.

II. BACKGROUND

A. Problem Definition

We consider a sequential decision process with multiple agents. We can define it as an n -agent stochastic game [1], [2] consisting of:

- set of agents $\mathcal{N} = \{1, \dots, n\}$,
- state space \mathcal{S} ,
- action space \mathcal{A}^i for each agent $i \in \mathcal{N}$,
- reward function \mathcal{R}^i for each agent $i \in \mathcal{N}$, defined as $\mathcal{R}_i : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$, where $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^n$,
- state transition probability function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$.

At each time step t , agents choose their actions a_i^t according to individual policies $\pi_i : \mathcal{S} \mapsto \Delta(\mathcal{A}^i)$ for state s^t . Based on the joint action $a^t = \{a_i^t\}_{i \in \mathcal{N}}$, state transits to s^{t+1} and each agent receives reward r_i^t . In independent learning, goal of each agent i is to find optimal policy π_i^* that maximizes expected return defined as

$$\mathbb{E}_{\pi_i, \pi_{-i}} \left[\sum_{t=0}^{\infty} \gamma^t r_i^t \right],$$

where $\pi_{-i} = \prod_{j \in \mathcal{N} \setminus \{i\}} \pi_j$ is joint (Cartesian product) policy of other agents, and $\gamma \in [0, 1)$ is a discount factor. This means, that optimal policy of agent i also depends on how the other agents act. Given this goal, the solution to a stochastic game will be (ϵ -)Nash equilibrium, where no agent can improve its performance by unilaterally changing its policy.

B. Independent Learning

In independent learning, each agent learns its policy π_i based solely on its own observations, actions and rewards, while completely ignoring the existence of other agents [1]. This approach essentially reduces a multi-agent problem to a series of single-agent problems, for which we can use single-agent reinforcement learning algorithms.

From the perspective of agent i , the policies π_j of other agents become part of the state transition function:

$$\mathcal{T}_i(s^{t+1}|s^t, a_i^t) \propto \sum_{a_{-i} \in \mathcal{A}^{-i}} \mathcal{T}(s^{t+1}|s^t, a_i^t, a_{-i}) \prod_{j \neq i} \pi_j(a_j|s^t),$$

where $-i$ means “all agents other than i .” During learning, each agent j will continue to update its policy π_j , and the action probabilities associated with π_j in each state s may change. Therefore, from the perspective of agent i , the transition function \mathcal{T}_i seems to be non-stationary. However, this perceived non-stationarity is due solely to the evolving policies π_j of the other agents. Consequently, independent learning approaches can lead to unstable learning and might not converge to a stable solution in the game.

We can alleviate this problem by allowing some centralized information. For example, during training, the algorithm can utilize the shared local information from all agents to update their policies. One of the most noticeable algorithms of *centralized training and decentralized execution* (CTDE) paradigm is actor-critic method known as MADDPG [3], which uses centralized action-value function as a critic. Access to global information helps coordinate the actions of agents, which can mitigate the effect of non-stationarity and contribute to more stable learning. However, centralization introduces issues with scalability, as the number of joint actions grows exponentially relative to the number of agents.

III. INTEGRATION OF ENVIRONMENT MODEL

A model of the environment refers to anything that an agent can utilize to predict the environment’s response to its actions [4]. Given a state and an action, the model provides a forecast of the resulting next state and reward. In the case of a stochastic model, there are many possible future states and rewards, each with probabilities assigned to their occurrence. These models can take two forms:

- *distribution model*, which provides the complete distribution of all possibilities and their probabilities,
- *sample model*, which generates a possible state and reward by sampling according to the assigned probabilities.

In both cases, model is used to simulate the environment and generate simulated experiences.

The term *planning* can have various meanings across different fields. In our context, planning refers to any computational process that uses a model as input to generate or refine a policy for interacting with the modeled environment. Planning can be integrated into reinforcement learning in two ways: 1) we can use the model to generate simulated experiences and use these experiences to improve the policy or value function (*background planning*); and 2) we can use the model to plan in predicted states in order to select actions (*decision-time planning*) [4].

Based on how learning is used, we can distinguish three main categories of planning-learning integration [5]:

- *model-based RL with a learned model*, where we learn both model and policy/value function (e.g. Dyna-Q [6]),
- *model-based RL with a known model*, where we plan over a known model, and only use learning for value function (e.g. AlphaZero [7]),
- *planning over a learned model*, where we learn model, and use it to plan locally, without learning policy or value function (e.g. Embed2Control [8]).

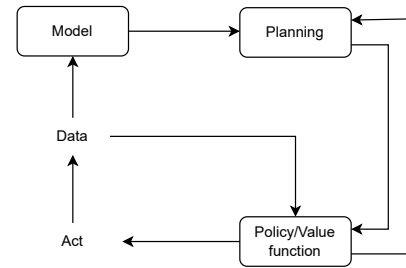


Fig. 1. Integration of planning and learning in Dyna architecture

Here, “planning over a learned model” may not be considered model-based RL, since no policy or value function is learned.

A. Dyna Architecture

The Dyna [9] architecture is an integrated framework that combines learning, planning, and reactive execution in the context of reinforcement learning (Fig. 1). Real experience is utilized by the planning agent to improve the model, making it more accurately reflect the real environment, as well as being used, similarly to model-free methods, for directly improving the value function and policy. The model is used to generate simulated experiences, which are then utilized to update the policy/function. Planning is done incrementally and can utilize world models often generated by learning processes, even if they are sometimes incorrect. If the model is accurate, planning significantly speeds up finding the optimal policy. In small tasks, it has been demonstrated to be true, even if the model also needs to be learned or if the environment changes [6].

Recent progress in Dyna-style MARL methods includes *Adaptive Opponent-wise Rollout Policy Optimization* (AORPO) [10]. AORPO explores ways to improve sample efficiency in stochastic games, where agents independently learn their policies but have the capacity to communicate with each other. It utilizes opponent models to generate model rollouts for a specific number of steps, determined by the validation error of the opponent model. Subsequent steps of the rollout involve requesting actions from the corresponding opponent through communication. Improving policies with data from environment models that predict joint agent actions helps with non-stationarity, but opponent models need full observability, which may not always be possible.

B. Heuristic Search

Classical decision-time planning methods are collectively known as *heuristic search* [4]. In heuristic search, for each encountered state, a large tree of possible future steps (based on the model) is considered. An approximate value function (typically designed by humans and never updated as a result of the search) is applied to the nodes at the ends of branches, and then these values are backed-up towards the current state, which is the root of the tree. When we eventually compute

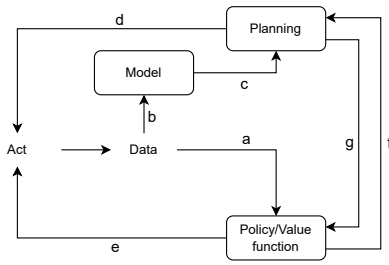


Fig. 2. Algorithmic connections between planning and learning

values of action nodes for the current state, the best of them is selected as current action. Greedy policies are like heuristic search but on a smaller scale. For example, to select a greedy action using the model and state value function, we need to predict future states for each possible action, consider the rewards and estimated values of those states, and then choose the action that yields the best outcome. Heuristic search can thus be understood as an extension of the greedy policy idea beyond a single step. Here, we assume we have a perfect model and an imperfect state value function. In such a scenario, deeper search usually leads to a better policy — if the search was carried out until the end of the episode, then indeed, the impact of the imperfect value function would be eliminated, and the action chosen in this way must be optimal. However, the deeper the search, the more computations are required.

An interesting example is the TD-Gammon program [11], created by Gerald Tesauro, which achieved a master-level performance in the game of backgammon. It uses a form of heuristic search to select moves and learns a value function over multiple self-play games. As a model, TD-Gammon utilized a priori knowledge about the probabilities of rolling specific dice outcomes and the assumption that the opponent always chose actions that TD-Gammon deemed best for it.

IV. DESIGNING THE ALGORITHM

In our experiments, we want to explore how the environment model can affect the issue of non-stationarity in independent learning. Specifically, we want to find out whether using a model (even if imperfect) helps the agent reduce the sense of environment non-stationarity caused by other learning agents.

According to [5], there are several possible algorithmic connections between planning and learning, as shown in Fig. 2: a) learning a policy/value function from real data, b) learning the model from real data, c) planning using the model, d) acting based on the outcome of planning, e) acting based on the policy/value function, f) using the policy to improve the planning procedure, g) using the planning result to update the policy/value function. We need to consider how to use the environment model in our algorithm.

Suppose we have an ideal environment model, which implicitly must be in line with the policies of the other agents, and a well-designed approximate value function. We could use heuristic search to choose optimal actions in each state.

Algorithm 1 Designed model-based RL algorithm

```

// Algorithm controls agent  $i$ 
1: Initialize predictive model  $\hat{T}_i$  and state-value function  $V^{\pi_i}$ 
2: Repeat for every episode:
3: for  $t = 0, 1, 2, \dots$  do
4:   Observe current state  $s^t$ 
5:   if explore with probability  $\varepsilon$  then
6:     Choose random action  $a_i^t \in \mathcal{A}_i$ 
7:   else
8:     for each action  $a_i \in \mathcal{A}_i$  do
9:       Obtain predicted next state  $\hat{s}^{t+1}$  and predicted
       reward  $\hat{r}_i^t$  from model  $\hat{T}_i$  using state  $s^t$  and action
        $a_i$ 
10:      Calculate value of the action  $AV(a_i) \leftarrow \hat{r}_i^t +$ 
        $\gamma V^{\pi_i}(\hat{s}^{t+1})$ 
11:    end for
12:    Choose action  $a_i^t = \arg \max_{a_i} AV(a_i)$ 
13:  end if
14:  (meanwhile, other agents  $j \neq i$  choose their actions  $a_j^t$ )
15:  Observe real reward  $r_i^t$  and next state  $s^{t+1}$ 
16:  Update predictive model  $\hat{T}_i(s^t, a_i^t)$  using  $s^{t+1}$  and  $r_i^t$ 
17:   $V^{\pi_i}(s^t) \leftarrow V^{\pi_i}(s^t) + \alpha[r_i^t + \gamma V^{\pi_i}(s^{t+1}) - V^{\pi_i}(s^t)]$ 
18: end for
    
```

A change in the environment dynamics (in this case, a change in someone’s policy) would not affect an agent that relies entirely on its own model (assuming the model is not updated). This way, we would get rid of the non-stationarity problem. However, the model wouldn’t be perfect anymore, so our future actions could be suboptimal. However, if the other agents also operate optimally, then there would be no change in dynamics in the first place.

In a situation with multiple learning agents, we cannot rely on a perfect environment model, and we must adjust it to accommodate the changing behavior of the other agents. Therefore, just like in the Dyna architecture, we need to update the model during learning. Given the inherent imperfection of the model, deep heuristic search is not feasible, as prediction errors in subsequent states will quickly accumulate. Therefore, we will plan only one step ahead.

We will approximate the values of future states using the state value function V^{π_i} for the current policy π_i of agent i (which is conditioned on the agent’s environment model). This means it can only be updated using on-policy data. To make updates, we will apply one-step TD learning using experiences from interactions with the real environment. Due to the imperfection of the model, using it to acquire simulated experiences to improve V^{π_i} may be infeasible.

We hope that this way of integrating planning and learning (also presented in Fig. 3) will, at least to some extent, reduce the problem of non-stationarity. Algorithm 1 presents pseudocode of the designed model-based RL method.

V. ARCHITECTURAL CHOICES

During the experiments, we use following hyperparameters:




Fig. 3. Integration of planning and learning in the designed algorithm

- learning rate $\alpha = 0.0001$,
- discount rate $\gamma = 0.9$.

To encourage exploration, we use the ε -greedy method. The parameter ε will initially be $\varepsilon_{start} = 0.99$, which will gradually decrease to $\varepsilon_{end} = 0.05$ according to the formula:

$$\varepsilon = \varepsilon_{end} + (\varepsilon_{start} - \varepsilon_{end})e^{(-k/1000)}$$

where k is the total number of steps taken by the agent since the beginning of the learning process.

A. Environment Model

The environment model is a sample model which, given an input observation, returns the predicted next observations for each possible action, the predicted rewards for taking those actions, and information on whether each of these actions leads to a terminal state. It is a simple feedforward neural network with parameters θ composed of 2 hidden layers, each with 64 neuron units using the ReLU activation function.

The output layer predicting the next observations has a size equal to the number of actions multiplied by the observation size. The output layer predicting the rewards has a width equal to the number of actions. Both of these output layers do not use nonlinear activation functions.

The output layer predicting whether the next state will be terminal has a size equal to the number of actions and uses the sigmoid activation function to represent the probability, where a value of 1 indicates that the model is certain it will be a terminal state.

In the implementation, we use a memory buffer that stores the history of the last 1000 steps of the agent. At each step, a randomly selected batch of 32 experience tuples $(s, a, r, s', final)$, is retrieved from the buffer, where $final$ is an indicator whether state s' is a terminal state. Subsequently, based on this batch, the parameters of the environment model of this agent are updated using the mean squared error (MSE) loss:

$$\mathcal{L}(\theta) = MSE(\hat{s}', s') + MSE(\hat{r}, r) + MSE(\hat{final}, final)$$

For parameter updates, we employ the AdamW [12] optimizer with a learning rate α .




Fig. 4. Simple Spread environment with 3 agents (blue circles represent agents, and black dots are destination landmarks)

B. State-Value Function

The state-value function V^{π_i} or agent i 's policy π_i is approximated using a feedforward neural network with parameters ϕ . This network consists of two hidden layers, similar to the environment model. It takes the state (observation) s as input, and its output layer returns a single value — the predicted state value $V^{\pi_i}(s)$ under policy π_i . The output layer does not use a nonlinear activation function. The network is updated at each timestep using the acquired tuple $(s, a, r, s', final)$ of real experience, utilizing the mean squared error loss function:

$$\mathcal{L}(\phi) = MSE(V^{\pi_i}(s), y_i)$$

where $y_i = r + \gamma V^{\pi_i}(s')$ if s' is not a terminal state, or $y_i = r$ otherwise. We use AdamW optimizer with a learning rate α for parameter updates.

VI. EXPERIMENTAL SETUP

A. Environment

Multi-Particle Environments (MPE) [3] is a collection of environments where agents (particles) are placed in a two-dimensional space with designated landmarks. They can interact with the environment and with each other to achieve specific goals that require cooperation or competition.

One particular environment of interest is the navigation with cooperation available in the *PettingZoo* [13] library under the name *Simple Spread*. Fig. 4 illustrates this game, with blue particles representing the agents. In this environment, agents cooperate physically (without communication) to reach multiple landmarks. They observe the positions of other agents and points of interest, and their reward depends on how close they are to these points. The goal is to cover all landmarks while avoiding collisions, for which they are penalized. Agents learn which landmark to head towards and navigate there while avoiding other agents. In our experiments the game terminates after 32 timesteps.

B. Comparison Algorithm

We want to compare our model-based RL algorithm with a model-free one. Since our designed algorithm resembles Q-learning, we will compare it with the well-known Deep Q-Network (DQN) algorithm [14].

In the implementation, we use experience memory to store the agent's experiences from the last 1000 time steps. The Q-network is similar to the state-value network in our algorithm, but instead of a single value, it outputs one value for each possible action. It also uses AdamW optimizer with learning rate α . The target network is updated in every step using soft update:

$$\theta_{target} \leftarrow \tau\theta_{current} + (1 - \tau)\theta_{target}$$

where $\tau = 0.005$, and θ_{target} and $\theta_{current}$ represent the parameters of the target network and the current network, respectively.

VII. TESTS AND RESULTS

We trained our implemented algorithm and the DQN algorithm in the Simple Spread environment with 3 agents for 150,000 episodes. We compared the performance of the DQN algorithm with that of our model-based algorithm, illustrating the learning trends over time. The rewards for both algorithms were tracked and plotted to observe their respective improvements in performance throughout the training process.

The plot in Fig. 5 presents a moving average of the total rewards over 1000 episodes for these two algorithms. As training progresses, the plot shows the smoothed trend of the rewards received by each algorithm, highlighting their performance improvements over time. The moving average helps in reducing the noise in the reward signal, providing a clearer picture of the algorithms' learning behaviors.

In this test, our model-based method achieves better asymptotic performance than model-free DQN. DQN is an off-policy algorithm (meaning it can use the experience generated by policies other than the behavior policy to train its value function, thus enabling the reuse of old experience, which can significantly speed up the learning process) and is known for its high sample efficiency in single-agent RL. The fact that our on-policy algorithm (contrary to off-policy, it cannot use the outer experience to improve value function) achieves similar sample efficiency in the scenario with three agents must, therefore, be attributed to its more effective handling of the non-stationarity problem.

A. Single Agent

To demonstrate that DQN is indeed more sample efficient in single-agent RL, we trained both algorithms in the same environment but with only a single agent. Fig. 6 shows their performance improvements during training (also using moving average over 1000 episodes). In this test, both methods achieve similar asymptotic performance, but DQN reaches it much faster.




Fig. 5. Comparison of DQN's and our model-based (MB) algorithm's learning performance in Simple Spread with 3 agents




Fig. 6. Comparison of DQN's and our model-based (MB) algorithm's learning performance in Simple Spread with 1 agent

B. More Agents

To determine if our model-based algorithm helps reducing non-stationarity, we compared it with DQN in Simple Spread environment with 4 and 6 agents. The presence of additional agents exacerbates the non-stationarity effect, allowing for a more discernible assessment of our algorithm's efficacy in addressing it. Fig. 7 and Fig. 8 show comparison plots for environments with 4 and 6 agents respectively. In both cases, the DQN agents failed to make progress in solving the problem due to the presence of non-stationarity, resulting in cyclic behavior. In contrast, our method showed signs of improved performance.

Conducted tests indicate that our model-based algorithm outperforms similar model-free method in multi-agent setups with decentralized learning. These setups are highly affected by non-stationarities, wherein our model-based approach demonstrates superior adaptability and performance.

VIII. CONCLUSION

We created a custom model-based algorithm for multi-agent reinforcement learning with decentralized training. Our algorithm uses the environment model to plan in predicted states in order to select actions. However we allow for the actions to be selected based on an imperfect model, enabling




Fig. 7. Comparison of DQN's and our model-based (MB) algorithm's learning performance in Simple Spread with 4 agent




Fig. 8. Comparison of DQN's and our model-based (MB) algorithm's learning performance in Simple Spread with 6 agent

the model to be learned and improved throughout the training process.

We tested our algorithm against the Deep Q-Network in a sample environment. As the number of agents increased, our model-based algorithm outperformed DQN. Therefore, it is reasonable to assume that this trend will continue with even more agents, although this needs to be confirmed with additional tests.

Our experimental findings suggest that utilizing an environment model can effectively help agents in mitigating the effects of non-stationarity induced by other adaptive agents. To confirm conclusively, additional experiments across a broader range of environments are required. Nevertheless, this work contributes to the understanding of how environment modeling can improve the effectiveness of multi-agent reinforcement learning algorithms. It presents an alternative approach to the challenges present in the field compared to most currently developed methods that rely on centralized training.

It may also be possible to use the model to generate simulated experience for improving the value function (as indicated by arrow g in Fig. 2) when the model is accurate enough (e.g., when model loss is sufficiently small). This could be the focus of further work.

REFERENCES

- [1] S. V. Albrecht, F. Christianos, and L. Schäfer, *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024. [Online]. Available: <https://www.marl-book.com>
- [2] L. S. Shapley, "Stochastic games*," *Proceedings of the National Academy of Sciences*, vol. 39, no. 10, pp. 1095–1100, 1953. doi: 10.1073/pnas.39.10.1095. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.39.10.1095>
- [3] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," *CoRR*, vol. abs/1706.02275, 2017. doi: 10.48550/arXiv.1706.02275. [Online]. Available: <http://arxiv.org/abs/1706.02275>
- [4] R. S. Sutton and A. G. Barto, *Reinforcement Learning*, 2nd ed., ser. Adaptive Computation and Machine Learning. Cambridge, MA: MIT Press, 2018. ISBN 978-0-262-03924-6. [Online]. Available: <http://incompleteideas.net/book/the-book.html>
- [5] T. M. Moerland, J. Broekens, A. Plaat, and C. M. Jonker, "Model-based reinforcement learning: A survey," 2022. doi: 10.48550/arXiv.2006.16712
- [6] R. S. Sutton, "Integrated architectures for learning, planning, and reacting based on approximating dynamic programming," in *Machine Learning Proceedings 1990*, B. Porter and R. Mooney, Eds. San Francisco (CA): Morgan Kaufmann, 1990, pp. 216–224. ISBN 978-1-55860-141-3. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9781558601413500304>
- [7] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. P. Lillicrap, K. Simonyan, and D. Hassabis, "Mastering chess and shogi by self-play with a general reinforcement learning algorithm," *CoRR*, vol. abs/1712.01815, 2017. doi: 10.48550/arXiv.1712.01815. [Online]. Available: <http://arxiv.org/abs/1712.01815>
- [8] M. Watter, J. T. Springenberg, J. Boedecker, and M. A. Riedmiller, "Embed to control: A locally linear latent dynamics model for control from raw images," *CoRR*, vol. abs/1506.07365, 2015. [Online]. Available: <http://arxiv.org/abs/1506.07365>
- [9] R. S. Sutton, "Dyna, an integrated architecture for learning, planning, and reacting," *SIGART Bull.*, vol. 2, no. 4, p. 160–163, jul 1991. doi: 10.1145/122344.122377. [Online]. Available: <https://doi.org/10.1145/122344.122377>
- [10] W. Zhang, X. Wang, J. Shen, and M. Zhou, "Model-based multi-agent policy optimization with adaptive opponent-wise rollouts," 2022. doi: 10.48550/arXiv.2105.03363
- [11] G. Tesauro, "Programming backgammon using self-teaching neural nets," *Artificial Intelligence*, vol. 134, no. 1, pp. 181–199, 2002. doi: 10.1016/S0004-3702(01)00110-2. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0004370201001102>
- [12] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 2019. doi: 10.48550/arXiv.1711.05101
- [13] J. K. Terry, B. Black, N. Grammel, M. Jayakumar, A. Hari, R. Sullivan, L. Santos, R. Perez, C. Horsch, C. Dieffendahl, N. L. Williams, Y. Lokesh, and P. Ravi, "Pettingzoo: Gym for multi-agent reinforcement learning," 2021.
- [14] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," 2013. doi: 10.48550/arXiv.1312.5602

Lower Bounds on Cardinality of Reducts for Decision Tables from Closed Classes

Azimkhon Ostonov

0000-0001-5763-9751

King Abdullah University of

Science and Technology (KAUST)

Thuwal 23955-6900, Saudi Arabia

Email: azimkhon.ostonov@kaust.edu.sa

Mikhail Moshkov

0000-0003-0085-9483

King Abdullah University of

Science and Technology (KAUST)

Thuwal 23955-6900, Saudi Arabia

Email: mikhail.moshkov@kaust.edu.sa

Abstract—In this research paper, we examine classes of decision tables that are closed under attribute (column) removal and changing of decisions associated with rows. For decision tables belonging to these closed classes, we investigate lower bounds on the minimum cardinality of reducts. Reducts are minimal sets of attributes that allow us to determine the decision attached to a given row. We assume that the number of rows in the decision tables from the closed class is not limited by a constant. We divide the set of these closed classes into two families. In one family, the minimum cardinality of reducts for decision tables is bounded by standard lower bounds of the form $\Omega(\log \text{cl}(T))$, where $\text{cl}(T)$ represents the number of decision classes in the table T . In the other family, these lower bounds can be significantly tightened to the form $\Omega(\text{cl}(T)^{1/q})$ for some natural number q .

I. INTRODUCTION

DECISION tables are a widely recognized method for organizing and presenting information that is crucial for decision-making. These tables have various applications in data analysis, such as classification problems, studying combinatorial optimization, fault diagnosis, and computational geometry, among others. They have been extensively studied and utilized in different fields, as evidenced by the works [1], [2], [3], [4], [5], [6], [7], [8], [9]. It is worth noting that finite information systems with a designated decision attribute, datasets with a selected class attribute, and partially defined Boolean functions, which are commonly used in various data analysis domains to represent decision problems, can all be naturally interpreted as decision tables.

In this study, we focus on classes of decision tables that exhibit closure properties regarding operations of attribute (column) removal and decision modification attached to rows. One of the most natural examples of such classes is the set of decision tables derived from information systems. This set forms a closed class of decision tables. However, the family of all closed classes of decision tables is more extensive than the family derived from information systems alone. For instance, the union of classes derived from two separate information systems is also a closed class. However, it is important to note that there may not exist a single information system from which this union can be derived as a closed class.

This work was supported by King Abdullah University of Science and Technology

We investigate lower bounds on the minimum cardinality of reducts for decision tables belonging to closed classes. Reducts are minimal sets of attributes that enable the recognition of the decision attached to a given row of the table. In rough set theory, reducts play a crucial role in feature selection, classification problem solving, and knowledge compression [2], [7], [10], [11], [12], [13], [14]. Therefore, determining the lower bounds on the minimum cardinality of reducts is of considerable significance in rough set theory.

In this study, we make the assumption that the number of rows in decision tables belonging to the closed class is not restricted by a constant. We categorize these closed classes into two families. In one family, the minimum cardinality of reducts for decision tables is bounded by standard lower bounds of the form $\Omega(\log \text{cl}(T))$, where $\text{cl}(T)$ represents the number of decision classes in the table T . In the other family, these lower bounds can be significantly tightened to the form $\Omega(\text{cl}(T)^{1/q})$ for some natural number q . The findings obtained from this research can be valuable for experts in the field of data analysis.

This paper is divided into six sections. Sections II and III provide the primary definitions and relevant results pertaining to decision tables and closed classes of decision tables. In Sect. IV, we explore lower bounds on the cardinality of reducts. Section V presents examples that are associated with closed classes of decision tables derived from information systems. Finally, Sect. VI summarizes the main findings and presents brief conclusions.

II. DECISION TABLES

Consider a nonempty finite set B with k elements, where $k \geq 2$. We define a B -decision table T as a rectangular table with n columns. Each column is labeled with attributes (specifically attribute names). The rows of the table consist of distinct tuples from B^n , and they are labeled with nonnegative integers representing decisions. In this context, $\text{Rows}(T)$ refers to the set of rows in table T , $N(T)$ represents the total number of rows in T , and $\text{cl}(T)$ represents the number of distinct decisions attached to the rows of T (also known as the number of decision classes in table T). The value n is referred to as the dimension of table T and is denoted as $\dim T$.

A test for a table T is defined as a set of attributes (columns) from the table T such that any two rows in T with different decisions have at least one differing attribute in the selected set of columns. A reduct for table T is a test for T where none of its proper subsets can serve as a test. The minimum number of attributes in a reduct for table T is denoted as $R(T)$. If the number of distinct decisions in T (i.e., $\text{cl}(T)$) is less than 2, then $R(T)$ is defined as 0.

Let $[T]$ represent the set of decision tables that can be derived from T using the following process: we are allowed to remove any number of attributes (columns) from T , retain only one row from each group of identical rows in the resulting table, and modify the decisions attached to the remaining rows in any desired manner.

A decision table T that has n columns is referred to as quasicomplete if there exist subsets B_1, \dots, B_n of the set B , each consisting of two elements, such that the Cartesian product $B_1 \times \dots \times B_n$ is a subset of $\text{Rows}(T)$. We use $I(T)$ to represent the highest dimension among all quasicomplete tables from $[T]$. The following statement immediately follows from Theorem 4.6 in the work [5].

Lemma 1. For arbitrary B -decision table T with $\text{cl}(T) \geq 2$,

$$N(T) \leq (k^2 \dim T)^{I(T)}.$$

III. CLOSED CLASSES OF DECISION TABLES

Consider a set C consisting of B -decision tables. We define C as a closed class of decision tables if it can be represented as the union of $[T]$ for all T belonging to C . In other words, $C = \bigcup_{T \in C} [T]$. A closed class C is referred to as nondegenerate if the number of rows in tables from C is not restricted by a constant upper bound.

Next, we introduce a parameter $I(C)$ for a nondegenerate closed class C of decision tables. If the parameter I is limited by a constant for all tables in class C , then we define $I(C)$ as the maximum value among all $I(T)$ for T in C . However, if there is no upper bound on I for the tables in class C , then we assign $I(C)$ the value of positive infinity.

Let's examine the characteristics of the function

$$N_C(n) = \max\{N(T) : T \in C, \dim T \leq n\}.$$

This function, defined over the set of natural numbers, represents the manner in which the number of rows in decision tables from the class C increases in the worst-case scenario as their dimension grows.

Lemma 2. Consider a nondegenerate closed class C of B -decision tables.

(a) If $I(C) < +\infty$, then $N_C(n) \leq (k^2 n)^{I(C)}$ for any natural n .

(b) If $I(C) = +\infty$, then $2^n \leq N_C(n) \leq k^n$ for any natural n .

Proof. (a) Suppose $I(C) < +\infty$. From Lemma 1, we can derive that $N_C(n) \leq (k^2 n)^{I(C)}$ for any natural n .

(b) Assume that $I(C) = +\infty$ and let n be a natural number. The inequality $N_C(n) \leq k^n$ is straightforward. Since

$I(C) = +\infty$, there exists a quasicomplete table $T_n \in C$ with a dimension $\dim T_n = n$. It is evident that $N(T_n) \geq 2^n$. Therefore, we have $2^n \leq N_C(n)$. \square

IV. BOUNDS ON CARDINALITY OF REDUCTS

To begin, we establish a preliminary statement.

Lemma 3. Consider a nondegenerate closed class C of B -decision tables, and let T be a decision table from C for which $\text{cl}(T) \geq 2$. Then

$$N_C(R(T)) \geq \text{cl}(T).$$

Proof. Assume that $R(T) = m$, and let $\{f_1, \dots, f_m\}$ be a reduct of table T with the smallest possible cardinality. We represent the table obtained by removing all attributes from T except for f_1, \dots, f_m as T' , where T' is a table from $[T]$. In this case, the number of rows in the table T' must be at least as large as the number of decision classes in T , which can be expressed as $N(T') \geq \text{cl}(T)$. Additionally, it is evident that $N(T') \leq N_C(m)$. Consequently, we can conclude that $N_C(m) \geq \text{cl}(T)$. \square

Theorem 1. Consider a nondegenerate closed class C of B -decision tables.

(a) If $I(C) < +\infty$, then $R(T) \geq \text{cl}(T)^{1/I(C)}/k^2$ for any table $T \in C$ with $\text{cl}(T) \geq 2$.

(b) If $I(C) = +\infty$, then $R(T) \geq \log_k \text{cl}(T)$ for any table $T \in C$ with $\text{cl}(T) \geq 2$.

(c) If $I(C) = +\infty$, the inequality $R(T) \geq \log_2 \text{cl}(T) + 1$ does not hold for infinitely many tables T from the class C where both the dimension (number of attributes) and the number of decision classes are not bounded from above by any fixed constants.

Proof. (a) Suppose $I(C) < +\infty$, $T \in C$, $\text{cl}(T) \geq 2$, and $R(T) = m$. Using Lemma 2, we can obtain that $N_C(m) \leq (k^2 m)^{I(C)}$. By Lemma 3, $N_C(m) \geq \text{cl}(T)$. Therefore, $(k^2 m)^{I(C)} \geq \text{cl}(T)$ and $m \geq \text{cl}(T)^{1/I(C)}/k^2$.

(b) Suppose $I(C) = +\infty$, $T \in C$, $\text{cl}(T) \geq 2$, and $R(T) = m$. Using Lemma 2, we can obtain that $N_C(m) \leq k^m$. By Lemma 3, $N_C(m) \geq \text{cl}(T)$. Therefore, $k^m \geq \text{cl}(T)$ and $m \geq \log_k \text{cl}(T)$.

(c) Consider a natural number n . Since $I(C) = +\infty$, there exists a quasicomplete decision table T_n in the class C with a dimension $\dim T_n = n$ and a number of decision classes $\text{cl}(T_n) \geq 2^n$. Let us assume that $R(T_n) \geq \log_2 \text{cl}(T_n) + 1$. Then we have $R(T_n) \geq \log_2 2^n + 1 = n + 1$. It is evident that $n \geq R(T_n)$. Therefore, the inequality $R(T_n) \geq \log_2 \text{cl}(T_n) + 1$ does not hold. \square

The statement (c) demonstrates that the bound mentioned in the statement (b) cannot be significantly improved.

V. CLOSED CLASSES OF DECISION TABLES DERIVED FROM INFORMATION SYSTEMS

The most common instances of closed classes of decision tables arise from infinite information systems. An infinite information system is defined as a triple $U = (A, F, B)$, where

A represents an infinite set of objects known as the universe, B is a finite set with k elements (where $k \geq 2$), and F is an infinite set of functions from A to B known as attributes. A problem within this context is specified by a finite number of attributes $f_1, \dots, f_n \in F$ where these attributes divide the universe A into nonempty domains, with each domain having fixed values for the attributes f_1, \dots, f_n . Each domain is associated with a decision. The objective is to determine the decision assigned to a given object $a \in A$ based on the domain to which a belongs.

A decision table represents this problem as follows: the table consists of n columns that are labeled with the attributes f_1, \dots, f_n . The rows of the table correspond to the domains and are labeled with the decisions assigned to those domains.

We use $\text{Tab}(U)$ to represent the set of decision tables that correspond to all problems over the information system U . It can be proven that $\text{Tab}(U)$ is a nondegenerate closed class of decision tables. We refer to this class as being derived from the information system U .

A subset $\{f_1, \dots, f_p\}$ of the set F is considered independent if there exist two-element subsets B_1, \dots, B_p of the set B such that for any tuple $(b_1, \dots, b_p) \in B_1 \times \dots \times B_p$, the system of equations

$$\{f_1(x) = b_1, \dots, f_p(x) = b_p\}$$

has a solution from A . If, for any natural number p , the set F contains an independent subset of cardinality p , then $I(\text{Tab}(U)) = +\infty$. Otherwise, $I(\text{Tab}(U))$ is equal to the maximum cardinality of an independent subset in the set F .

Next, we examine some examples of infinite information systems provided in the book [6].

Example 1. Consider the Euclidean plane P and a straight line l within it. This line divides the plane into two open half-planes, denoted as h_1 and h_2 , along with the line l itself. We assign an attribute to the line l , where this attribute takes the value 0 for points in h_1 and the value 1 for points in h_2 and on the line l . We denote the set of attributes corresponding to all lines in P as F_P , and we define the information system $U_P = (P, F_P, \{0, 1\})$.

In this system, there exist two lines that divide the plane P into four domains. However, there are no three lines that can divide P into eight domains. As a result, we have $I(\text{Tab}(U_P)) = 2$. For any table $T \in \text{Tab}(U_P)$ with the number of distinct decisions in the table $\text{cl}(T) \geq 2$, we have a lower bound on the minimum cardinality of reducts $R(T)$ given by $R(T) \geq \text{cl}(T)^{1/2}/4$. This lower bound is significantly tighter than the standard bound $R(T) \geq \log_2 \text{cl}(T)$.

Example 2. Consider two natural numbers, m and t . We use $\text{Pol}(m)$ to represent the set of polynomials with integer coefficients that depend on variables x_1, \dots, x_m . Similarly, $\text{Pol}(m, t)$ refers to the set of polynomials from $\text{Pol}(m)$ that have a degree no greater than t . We define information systems $U(m)$ and $U(m, t)$ in the following way: $U(m) = (\mathbb{R}^m, F(m), E)$ and $U(m, t) = (\mathbb{R}^m, F(m, t), E)$, where \mathbb{R} is the set of real numbers, $E = \{-1, 0, +1\}$, $F(m) = \{\text{sign}(p) :$

$p \in \text{Pol}(m)\}$, $F(m, t) = \{\text{sign}(p) : p \in \text{Pol}(m, t)\}$, and $\text{sign}(x) = -1$ if $x < 0$, $\text{sign}(x) = 0$ if $x = 0$, and $\text{sign}(x) = +1$ if $x > 0$. It can be demonstrated that $I(\text{Tab}(U(m))) = +\infty$ and $I(\text{Tab}(U(m, t))) < +\infty$. Consequently, for any natural number m and any table T from $\text{Tab}(U(m))$ such that $\text{cl}(T) \geq 2$, we have a lower bound for the minimum cardinality of reducts $R(T)$ given by $R(T) \geq \log_3 \text{cl}(T)$. This bound cannot be significantly improved.

Similarly, for any natural numbers m and t , and any table T from $\text{Tab}(U(m, t))$ such that $\text{cl}(T) \geq 2$, we have a lower bound for the minimum cardinality of reducts $R(T)$ given by $R(T) \geq \text{cl}(T)^{1/q}/9$ for some natural number q .

VI. CONCLUSIONS

This research paper introduces a division of nondegenerate closed classes of decision tables into two distinct families. In one family of closed classes, the minimum cardinality of reducts for decision tables is bounded by standard lower bounds, specifically $\Omega(\log \text{cl}(T))$, where $\text{cl}(T)$ represents the number of decision classes in the table T . In the other family of closed classes, these lower bounds can be significantly improved, reaching the form of $\Omega(\text{cl}(T)^{1/q})$ for some natural number q .

Acknowledgements

Research reported in this publication was supported by King Abdullah University of Science and Technology (KAUST).

REFERENCES

- [1] E. Boros, P. L. Hammer, T. Ibaraki, and A. Kogan, "Logical analysis of numerical data," *Math. Program.*, vol. 79, pp. 163–190, 1997.
- [2] I. Chikalov, V. V. Lozin, I. Lozina, M. Moshkov, H. S. Nguyen, A. Skowron, and B. Zielosko, *Three Approaches to Data Analysis - Test Theory, Rough Sets and Logical Analysis of Data*, ser. Intelligent Systems Reference Library. Springer, 2013, vol. 41.
- [3] J. Fürnkranz, D. Gamberger, and N. Lavrac, *Foundations of Rule Learning*, ser. Cognitive Technologies. Springer, 2012.
- [4] E. Humby, *Programs from Decision Tables*, ser. Computer Monographs. Macdonald, London and American Elsevier, New York, 1973, vol. 19.
- [5] M. Moshkov, "Time complexity of decision trees," in *Trans. Rough Sets III*, ser. Lecture Notes in Computer Science, J. F. Peters and A. Skowron, Eds., Springer, 2005, vol. 3400, pp. 244–459.
- [6] M. Moshkov and B. Zielosko, *Combinatorial Machine Learning - A Rough Set Approach*, ser. Studies in Computational Intelligence. Springer, 2011, vol. 360.
- [7] Z. Pawlak, *Rough Sets - Theoretical Aspects of Reasoning about Data*, ser. Theory and Decision Library: Series D. Kluwer, 1991, vol. 9.
- [8] S. L. Pollack, H. T. Hicks, and W. J. Harrison, *Decision Tables: Theory and Practice*. John Wiley & Sons, 1971.
- [9] L. Rokach and O. Maimon, *Data Mining with Decision Trees - Theory and Applications*, ser. Series in Machine Perception and Artificial Intelligence. World Scientific, 2007, vol. 69.
- [10] Z. Pawlak and A. Skowron, "Rudiments of rough sets," *Inf. Sci.*, vol. 177, no. 1, pp. 3–27, 2007.
- [11] D. Slezak, "Approximate entropy reducts," *Fundam. Informaticae*, vol. 53, no. 3–4, pp. 365–390, 2002.
- [12] S. Stawicki, D. Slezak, A. Janusz, and S. Widz, "Decision bireducts and decision reducts - a comparison," *Int. J. Approx. Reason.*, vol. 84, pp. 75–109, 2017.

- [13] A. Janusz and S. Stawicki, "Reducts in rough sets: Algorithmic insights, open source libraries and applications (tutorial – extended abstract)," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, ser. *Annals of Computer Science and Information Systems*, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 30. IEEE, 2022. doi: 10.15439/2022F261 p. 279–288. [Online]. Available: <http://dx.doi.org/10.15439/2022F261>
- [14] B. K. Vo and H. S. Nguyen, "Feature selection and ranking method based on intuitionistic fuzzy matrix and rough sets," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems*, ser. *Annals of Computer Science and Information Systems*, M. Ganzha, L. Maciaszek, M. Paprzycki, and D. Ślęzak, Eds., vol. 35. IEEE, 2023. doi: 10.15439/2023F0002 p. 71–71. [Online]. Available: <http://dx.doi.org/10.15439/2023F0002>

Automatic Generation of OpenCL Code through Polyhedral Compilation with LLM

Marek Palkowski, Mateusz Gruzewski
West Pomeranian University of Technology in Szczecin
ul. Zolnierska 49, 71-210 Szczecin, Poland
Email: mpalkowski@zut.edu.pl

Abstract—In recent years, a multitude of AI solutions has emerged to facilitate code generation, commonly known as Language Model-based Programming (LLM). These tools empower programmers to automate their work. Automatic programming also falls within the domain of optimizing compilers, primarily based on the polyhedral model, which processes loop nests concentrating most computations. This article focuses on harnessing LLM tools to generate OpenCL code for non-serial polyadic dynamic programming kernels.[1] We have chosen the Nussinov RNA folding computational task, previously employed to test polyhedral compilers in optimizing kernels with non-uniform dependences. The code generated in OpenMP by polyhedral optimizers is limited to CPU computations. We automatically convert it into the OpenCL standard using ChatGPT-3.5 through its source-to-source queries to extend the number of possible platforms. The validity and efficiency of the generated code were verified on various CPUs and GPUs from different manufacturers.

I. INTRODUCTION

CODE generation using LLM (Language Model-based Programming) tools has garnered significant interest among programmers and researchers in recent times. This represents a novel form of automatic programming. Generating code effortlessly is also within the domain of polyhedral compilers. Source-to-source techniques enable the generation of parallel and localized code through stages of syntax and loop dependency analysis, transformations, and loop generation. Thus, it is evident that leveraging LLM models for High-Performance Computing (HPC) opens up new possibilities.

Using the capabilities of the ChatGPT tool, we aim to automatically generate OpenCL code [2] based on existing OpenMP code [3] optimized with polyhedral tools. This allows us to execute the code on any graphics card that supports OpenCL. The test code will be based on the Nussinov algorithm [4] for RNA prediction with program loop nets containing many non-uniform loops and posing a challenge for polyhedral tools.

The Nussinov algorithm is a non-serial polyadic dynamic programming (NPDP) kernel, used to assess the efficiency of tiled code generated by advanced optimizing compilers [5], [6], [7], [8]. NPDP dependence patterns, the most complex category of Dynamic Programming (DP), exhibit non-uniform dependences characterized by irregularities and expressed using affine expressions. Dynamic Programming involves finding optimal solutions for simpler instances of a problem and extending them to solve larger instances. Additional difficulty

arises from parallelism with synchronization. The goal for GPT is to generate a sequential loop spawning a kernel for the GPU in the OpenCL standard.

The remaining sections of the paper are organized as follows. The subsequent section provides a concise overview of Language Model-based Programming (LLM) tools for High-Performance Computing (HPC) approaches. Section three provides a brief explanation of the Nussinov algorithm, while the fourth section delves into polyhedral optimization for this kernel. The following section introduces the process of generating OpenCL code using ChatGPT, utilizing OpenMP codes from Traco [9], Dapt [10], and Pluto [11] for the Nussinov loop nests. The experimental study assesses the efficiency of the generated codes for CPUs and GPUs. Finally, the last section concludes the paper and outlines potential avenues for future work.

II. RELATED WORK

Polyhedral frameworks for loop transformation comprise three essential stages: dependency tests, loop transformation via affine operations on polytopes, and code generation from modified loop polytopes. The central process of loop transformation varies across established compilers like Traco, Pluto, and Dapt, yet it plays a pivotal role in optimizing loop structures within program code.

Loop tiling, also referred to as loop blocking or loop partitioning, stands as a compiler optimization technique aimed at enhancing cache utilization and bolstering the performance of loop-based computations [12]. This method involves breaking down a loop into smaller blocks, or tiles, which can be efficiently accommodated within the cache memory. By organizing data access in close proximity to memory, loop tiling diminishes cache misses and optimizes memory access patterns, thereby contributing to overall performance improvements. Additionally, these tiles facilitate parallel processing, further amplifying computational efficiency.

However, achieving parallelization of NPDP kernels often involves employing established strategies such as loop-skewing. The form of multi-threaded code within polyhedral compilers hinges on the tiling algorithm adopted.

The PLUTO compiler [11] leverages the affine transformation framework (ATF) to produce parallel tiled code, utilizing loop transformations to bolster multi-threading capabilities and enhance data locality. It optimizes tiling hyperplanes by

employing an embedded Integer Linear Programming (ILP) cost function, thereby achieving efficient parallelism while minimizing communication overhead in the processor space. TRACO [9], on the other hand, utilizes the transitive closure of dependence relation graphs to generate valid target tiles, rectifying them by eliminating invalid dependence destinations. DAPT [10] addresses non-uniform dependencies by approximating them to uniform counterparts, thereby simplifying complexities associated with nonlinear time-tiling constraints. Unlike PLUTO, DAPT and TRACO support three-dimensional tiling and benchmarks like nussinov, nw, and sw. It's worth noting that Pluto faces limitations in parallelizing mcc code [13].

A drawback of these solutions is the inability to generate code for graphics cards, significantly limiting their applicability on platforms other than CPUs. There have been compilers and converters designed for heterogeneous computing, adhering to standards such as OpenMP and CUDA, for many years, such as Par4All [14] or Cetus [15], but they lack an optimizing engine at the level of the polyhedral model or rely on basic transformations.

The PPCG compiler [16], developed a decade ago, excels in producing optimized GPU code through polyhedral optimization and is actively maintained. Leveraging the ATF framework (specifically, the Pluto algorithm) and the ISL library, this tool generates CUDA [17] and OpenCL codes, showcasing successful performance in stencils and Polybench kernels. However, when applied to NPDP codes, several challenges emerged. PPCG 0.9.1 lacks the capability to generate parameterized code, necessitating constant parameter values during compilation. Consequently, the resulting code contains numerous constants tied to specific parameter values, demanding code regeneration for each parameter set. Hence, we opted to exclude the PPCG compiler from our experimental study.

The evaluation in papers [18], [19], [20] indicates that OpenMP and OpenACC compilers demonstrate proficiency in generating efficient parallel code for simpler cases with GPUs. However, as the complexity of the code increases, a notable performance gap emerges between CUDA or OpenCL and OpenACC, OpenMP. Specifically, when examining memory access patterns such as sum reduction, OpenMP exhibits significantly slower performance compared to the same optimized reduction pattern implemented in CUDA or OpenCL.

Alternatively, developers have the option to utilize solutions like LLM for creating a prototype of the target code, evaluating its performance, and subsequently integrating it into tool implementations. Nichols et al. [1] have further refined the application of LLMs to improve the generation of OpenMP pragmas, with extensions to MPI cases. In a related context, Chen et al. [21] introduced LM4HPC, a framework specifically tailored for HPC tasks using LLMs. They tackled the challenge of limited training and evaluation datasets in HPC by proposing an approach to identify parallelism in code through machine learning techniques.

In the two recent studies, the topic of generating efficient code for classical benchmarks using artificial intelligence (AI)

has garnered attention in the scientific literature. Godoy et al. [22] delved into AI-driven generative capabilities, focusing on fundamental numerical kernels in high-performance computing (HPC). Their evaluation encompassed various programming models like OpenMP and CUDA, across languages such as C++ or Python, utilizing both CPU and GPU processing. GitHub Copilot [23], powered by OpenAI Codex [24], was employed to generate multiple implementations based on prompt variations. Subsequently, Pedro et al. [25] revisited the experimental investigation using the Llama-2 engine [26], aiming to generate high-quality HPC codes for the same benchmarks and programming language models. Despite Llama-2's focus on providing optimized code solutions, the study noted a trade-off in terms of reliability when compared to Copilot.

Hence, however, while OpenCL demands significantly more programming effort, applying LLM techniques may render the code generation process easier and improve performance.

III. NUSSINOV RNA FOLDING

Nussinov pioneered one of the earliest attempts at computationally efficient RNA folding using the base pair maximization approach in 1978 [4]. An RNA sequence comprises a chain of nucleotides from the alphabet G (guanine), A (adenine), U (uracil), C (cytosine). The Nussinov algorithm addresses the challenge of predicting RNA non-crossing secondary structures by calculating the maximum number of base pairs for sub-sequences. It initiates the process with sub-sequences of length 1 and incrementally builds upwards, storing the result of each sub-sequence in a DP array.

Let N be a $n \times n$ Nussinov matrix and $\sigma(i, j)$ be a function which returns 1 if RNA[i], RNA[j] are a pair in the set (AU, UG, GC) and $i < j - 1$, or 0 otherwise. Then the following recursion $N(i, j)$ is defined over the region $1 \leq i \leq j \leq n$ as

$$N_{i,j} = \max(N_{i+1,j-1} + \sigma(i, j), \max_{1 \leq k \leq n} (N_{i,k} + N_{k+1,j})) \quad (1)$$

and zero elsewhere [27].

The equation leads directly to the C/C++ code with triple-nested loops presented in Listing 1 [5].

Listing 1. Nussinov loop nest.

```

for (i = N-1; i >= 0; i--) {
  for (j = i+1; j < N; j++) {
    for (k = 0; k < j-i; k++) {
      S[i][j] = MAX(S[i][k+i] + S[k+i+1][j], S[i][j]);
    }
    S[i][j] = MAX(S[i][j], S[i+1][j-1] + signa(i, j));
  }
}

```

IV. POLYHEDRAL OPTIMIZATION FOR NUSSINOV LOOP NESTS

The polyhedral model represents loop nests as polyhedra with affine loop bounds and schedules. This model provides a foundation for advanced loop transformations and the analysis of data dependences. By harnessing the power of the polyhedral model, compilers can automatically optimize loops,

enhance performance (especially in terms of locality with loop tiling), and exploit parallelism (particularly with loop skewing for NPDP codes) [28].

The polyhedral model is implemented in compilers such as Pluto [29], Dapt [10], and Traco [30], each based on ATF, space-time tiling, and tile correction, respectively. Pluto excels in generating well-balanced affine schedules, but for NPDP codes, the framework can address a set of affine equations to tile all loop nests [30]. Traco compiler generates 3D tiles using the transitive closure of the dependence graph of the union of loop dependences. Further transformations for 3D-tiling of NPDP codes were implemented in the Dapt compiler by dividing the iteration space into timed parallel spaces. Dapt addresses irregularities in code obtained in Traco, providing a comprehensive solution to optimize and refine the generated code [31].

Listing 2. The OpenMP Pluto code for Nussinov RNA folding code.

```

for (t2=1; t2<=N-1; t2++) {
  lbp=t2; ubp=N-1;
#pragma omp parallel for private(lbv, ubv, t4, t6)
  for (t4=lbp; t4<=ubp; t4++) {
    for (t6=0; t6<=t2-1; t6++) {
      S[-t2+t4][t4] = MAX(S[-t2+t4][t6+(-t2+t4)]
        + S[t6+(-t2+t4)+1][t4], S[-t2+t4][t4]);
    }
    S[-t2+t4][t4] = MAX(S[-t2+t4][t4],
      S[-t2+t4+1][t4-1] + pair((-t2+t4), t4));
  }
}

```

To parallelize the Nussinov RNA folding code, compilers apply the well-known loop skewing transformation [30]. This transformation yields code in which the outermost loop becomes serial, enabling the parallelization of the remaining loops. Listing 2 showcases the parallel code generated using Pluto for the Nussinov code presented in Listing 1. This code, along with the equivalents from Traco and Dapt, serves as input for the OpenCL code generator in ChatGPT.

VI. OPENCL CODE GENERATION WITH GPT-3.5

To generate codes, we utilize Traco’s output code adhering to OpenMP standards. GPT serializes the code; however, if we do not specify which program loop nest is parallel by placing atomic functions, we provide hints about parallel loop nests, and then GPT correctly prepares kernel spawning. The remainder of the code is generated automatically, including context and kernel source loading, as well as memory transfer. The main skeleton of code is presented on Listing 3. Finally, we instruct GPT to generate separate kernels for Pluto and Dapt in a similar manner as for TRACO. The kernel of TRACO code is presented on Listing 4.

GPT adeptly prepares source kernels and seamlessly integrates them into the main code. It facilitates memory allocation, platform management, context selection, and command queue building. This tool streamlines program and kernel argument construction. It effectively generates a serial outermost loop that spawns parallel kernels on the GPU. GPT also auto-generates code for time measurements and compares host and device output arrays. The inclusion of OpenCL object releases is automated at the end of the program.

In kernels, only 1-dimensional arrays are accepted. ChatGPT-3.5 linearizes 2-dimensional arrays and appropriately sets loop bounds in the target OpenCL kernels for each polyhedral optimizer—Traco, Pluto, and Dapt — within the generated OpenMP codes.

The full documentation of the GPT session is available on the GitHub repository page: <https://markpal.github.io/fedcsis24/>, accessed on 1 March 2024. We employed straightforward, communicative English, which proved sufficient for GPT.

VI. EXPERIMENTAL STUDY

In the experimental study, our primary objective was to validate the correctness of the codes generated by the language model. Unlike polyhedral compilers, which often rely on mathematical proofs, our approach necessitated empirical testing due to the absence of such proofs. We began by scrutinizing the structure of the generated codes, focusing on aspects like loop spawning in GPU kernels, kernel arguments, and linearized array addresses. Furthermore, we compared the values of output arrays with those obtained from computations performed on the host CPU using OpenMP, generated with Pluto, Traco and Dapt. Remarkably, the results exhibited consistency between the outputs. While GPT occasionally produced trivial errors, these were easily rectified, yielding code that met the required standards.

To evaluate the performance of the generated OpenCL codes for the studied Nussinov kernel, we conducted experiments on two modern Intel processors and three graphic cards from NVIDIA, AMD and Intel. Our assessment utilized an Intel Xeon Gold 6326 machine, featuring 36 hardware threads running at 3.5 GHz in turbo mode, alongside a substantial 48 MB L3 cache and 128 GB of RAM. Complementing the CPU, the system incorporated an NVIDIA A100 Tensor Core GPU equipped with 6912 CUDA cores and 80 GB of memory. We tested also the second card, AMD Radeon RX 6700S. It uses the Navi 22 chip based on the new RDNA 2 architecture. The 128 Bit memory system connects 8 GB GDDR6 with 2 GHz memory clock. Furthermore, the RX6800S includes 32 MB Infinity Cache. We analyse also the Intel Core i5 12th 1235-U CPU with 12MB Cache L3 Cache and 16MB RAM with Intel Iris Xe GPU 1.4GHz with 80 execution units.

The experimental setup operated on the Ubuntu 22.04 operating system, and we compiled the programs using the Intel C Compiler (icc 2021), gcc 11.4.0, and clang 14.0 for OpenCL with the -O3 optimization flag. The codes utilized in our study are accessible via the repository link: <https://github.com/markpal/fedcsis24/>, accessed on 8 April 2024.

The Table 1 contains measurements of the execution time for the Xeon Gold processor and the A100 and Radeon graphics cards for the original code and OpenMP tiled code for Traco, Dapt, Pluto compiled with the Intel C++ compiler, and OpenCL compiled with Clang. OpenCL enables the utilization of polyhedral codes for graphics cards. The execution times for the Tesla A100 are significantly faster than the tiled code on the CPU, and while AMD codes are slower, they still show

Listing 3. A wide listing float, single column

```

#include <CL/cl.h>
...
int main() {
... // declarations, source of kernel load
int n = 30000;    int* h_S, *cpu_S;    char* kernelSource = (char*)malloc(fileSize + 1);
FILE* file = fopen("computeS.cl", "r"); fread(kernelSource, 1, fileSize, file);
... // memory allocation
h_S = (int*)malloc(n * n * sizeof(int));
cpu_S = (int*)malloc(n * n * sizeof(int));
... // OpenCL classes for target platform
cl_int err;    cl_platform_id cpPlatform[2];    cl_uint platf_num;
cl_device_id device;    err = clGetPlatformIDs(1, cpPlatform, &platf_num);
...// Context, Queue and Program Build
cl_context context = clCreateContext(0, 1, &device, NULL, NULL, &err);
cl_command_queue queue = clCreateCommandQueue(context, device, 0, &err);
cl_program program = clCreateProgramWithSource(context, 1, sources, NULL, &err);
clBuildProgram(program, 1, &device, NULL, NULL, NULL);

cl_int build_status;
clGetProgramBuildInfo(program, device, CL_PROGRAM_BUILD_STATUS, sizeof(cl_int), &build_status, NULL);

if (build_status != CL_SUCCESS) {
    // Print compilation errors
...    return 1; // or some other error code
}
// Kernel & buffer init, S array passing
cl_kernel kernel = clCreateKernel(program, "computeS_pluto", &err);
cl_mem d_S = clCreateBuffer(context, CL_MEM_READ_WRITE, n * n * sizeof(int), NULL, &err);
clEnqueueWriteBuffer(queue, d_S, CL_TRUE, 0, n * n * sizeof(int), h_S, 0, NULL, NULL);
... // Arguments to kernel
clSetKernelArg(kernel, 0, sizeof(cl_mem), &d_S);
clSetKernelArg(kernel, 1, sizeof(int), &n);
clSetKernelArg(kernel, 2, sizeof(int), &chunk);
... // Nussionov calculation on device
auto gpu_start = std::chrono::high_resolution_clock::now();
for (int cl = 1; cl < 2 * n - 2; cl += 1) {
    clSetKernelArg(kernel, 3, sizeof(int), &cl);
    clEnqueueNDRangeKernel(queue, kernel, 1, NULL, &globalSize, NULL, 0, NULL, NULL);
    clFinish(queue);
}
auto gpu_end = std::chrono::high_resolution_clock::now();
// Array S receive
clEnqueueReadBuffer(queue, d_S, CL_TRUE, 0, n * n * sizeof(int), h_S, 0, NULL, NULL);
// Host computation for OpenCL code validation
auto cpu_start = std::chrono::high_resolution_clock::now();
// host computation ...
auto cpu_end = std::chrono::high_resolution_clock::now();
for (int i = 0; i < n * n; i++) // host validation
    assert(h_S[i] == cpu_S[i]);
... // OpenCL object releases
clReleaseMemObject(d_S);    clReleaseKernel(kernel);    clReleaseProgram(program);
clReleaseCommandQueue(queue);    clReleaseContext(context);
...
}

```

acceleration compared to the original code using the same code as the NVIDIA card. The measurements were conducted for RNA sequences with nucleotide counts ranging from 1000 to 30000.

In the subsequent Table 2, the timing results for Intel hardware; i5 processor with dedicated Iris Xe graphics are compared. Tiled code was compiled with the gcc compiler, and GPU code was compiled with Clang. Although for Dapt codes, which feature an aggressive tiling algorithm for NPDP codes, the Iris Xe sometimes lags behind, for TRACO tile-corrected codes, the graphics card exhibits comparable performance, and even faster for Pluto codes. However, Pluto codes do not tile

the innermost loop [30].

In summary, correct and efficient parallel code was generated on graphics cards. Although the code is the same for graphics cards, three different kernels constructed using LLM based on the output OpenMP codes from Pluto, Traco, and Dapt were placed in the OpenCL kernel file. GPU of each manufacturer showed acceleration compared to sequential code executed on the host.

VII. CONCLUSION

In this study, we successfully generated efficient and correct OpenCL code, which we verified using cards from three

Listing 4. A wide listing float, single column

```

__kernel void computeS_pluto(__global int* d_S, int n, int CHUNK_SIZE, int t2) {
    int globalThreadId = get_global_id(0);
    int t4_base = globalThreadId * CHUNK_SIZE + t2;
    for (int offset = 0; offset < CHUNK_SIZE && (t4_base + offset) <= n - 1; offset++) {
        int t4 = t4_base + offset;
        for (int t6 = 0; t6 <= t2 - 1; t6++) {
            d_S[(-t2 + t4) * n + t4] = max(d_S[(-t2 + t4) * n + t6 + (-t2 + t4)]
                + d_S[(t6 + (-t2 + t4) + 1) * n + t4], d_S[(-t2 + t4) * n + t4]);
        }
        d_S[(-t2 + t4) * n + t4] = max(d_S[(-t2 + t4) * n + t4], d_S[(-t2 + t4 + 1) * n + t4 - 1]
            + pair(-t2+t4, t4));
    }
}

```

TABLE I
TIME EXECUTION IN SECONDS FOR XEON GOLD, NVIDIA A100 AND AMD RADEON, AND OPENMP/OPENCL LIBRARIES.

size	XEON GOLD				NVIDIA A100			AMD Radeon		
	original	traco	dapt	pluto	dapt	traco	pluto	dapt	traco	pluto
	icc	icc + openmp			clang + opencl			clang + opencl		
1000	0,35	0,23	0,18	0,05	0,24	0,23	0,13	0,28	0,18	0,28
2500	7,94	1,68	0,46	1,25	1,48	1,47	0,76	1,68	1,07	1,68
5000	165,32	7,84	6,71	4,33	8,18	8,15	4,24	7,23	5,14	7,23
7500	754,11	21,29	21,09	16,46	44,59	43,16	16,98	23,12	21,76	22,18
10000	1696,16	44,83	48,49	115,71	55,47	56,62	28,37	67,19	64,11	67,06
15000	5183,39	133,99	127,81	398,54	105,76	106,2	56,6	179,80	193,37	181,90
20000	13924,45	295,48	283,14	1100,78	243,53	239,98	119,8	582,49	602,83	580,06
30000	60565,98	972,52	917,96	4056,29	494,24	496,68	307,1	2 142,12	2 267,41	2 197,76

TABLE II
TIME EXECUTION IN SECONDS FOR CORE I5 AND IRIS XE, AND OPENMP/OPENCL LIBRARIES.

size	Intel Core i5				Intel Iris Xe		
	original	dapt	traco	pluto	dapt	traco	pluto
	gcc	gcc + openmp			clang + opencl		
1000	0,46	0,28	0,38	0,12	0,64	0,64	0,42
2500	14,15	2,35	3,22	2,08	3,57	3,84	2,67
5000	114,06	12,24	19,89	21,5	19,3	19,24	16,9
7500	404,86	40,2	61,86	101,9	65,76	66,06	57,27
10000	2261,75	114,27	181,33	717,31	190,89	191,59	161,85
15000	>3000	368,08	551,48	2438	437,53	442,03	435,28
20000	>3000	851,1	1478,94	>3000	1591,09	1589,77	1286,72

different manufacturers. This expanded the possibilities of utilizing results from polyhedral compilers and the concept of automated programming itself. The OpenCL code was created without writing a single line of code.

In future studies, we intend to explore further possibilities of LLM tools, including GPT-4, Github Copilot, or Llama-2, for various NPDP benchmarks. We are interested in the potential of utilizing once constructed code on other similar codes with non-uniform dependencies. Additionally, we plan to investigate optimization techniques, such as the use of shared memory, to enhance the performance of the generated code. LLM tools expand the capabilities of automated programming and appear to be a promising solution in heterogeneous programming.

REFERENCES

- [1] D. Nichols, A. Marathe, H. Menon, T. Gamblin, and A. Bhatele, "Modeling parallel programs using large language models," 2023, accessed on: 2024-01-11.
- [2] Khronos OpenCL Working Group, *The OpenCL Specification, Version 1.1*, 2011. [Online]. Available: <https://www.khronos.org/registry/cl/specs/opencl-1.1.pdf>
- [3] OpenMP Architecture Review Board, "OpenMP application program interface version 5.2," <https://www.openmp.org/specifications>, 2021, accessed on: 2023-10-22.
- [4] R. Nussinov *et al.*, "Algorithms for loop matchings," *SIAM Journal on Applied mathematics*, vol. 35, no. 1, pp. 68–82, 1978.
- [5] R. T. Mullapudi and U. Bondhugula, "Tiling for dynamic scheduling," in *Proceedings of the 4th International Workshop on Polyhedral Compilation Techniques*, S. Rajopadhye and S. Verdoolaege, Eds., Vienna, Austria, Jan. 2014.
- [6] D. Wonnacott, T. Jin, and A. Lake, "Automatic tiling of "mostly-tileable" loop nests," in *5th International Workshop on Polyhedral Compilation Techniques*, Amsterdam, 2015.
- [7] R. Chowdhury, , and *et al.*, "Autogen: Automatic discovery of efficient recursive divide-8-conquer algorithms for solving dynamic programming problems," *ACM Transactions on Parallel Computing*, vol. 4, no. 1, pp. 1–30, oct 2017. doi: 10.1145/3125632
- [8] W. Bielecki, P. Blaszyński, and M. Poliwoła, "3d parallel tiled code implementing a modified Knuth's optimal binary search tree algorithm," *Journal of Computational Science*, vol. 48, p. 101246, jan 2021. doi: 10.1016/j.jocs.2020.101246
- [9] W. Bielecki and M. Palkowski, "A parallelizing and optimizing compiler - traco," <http://traco.sourceforge.net>, 2013, accessed on: 2024-01-11.

- [10] W. Bielecki and M. Poliwoda, "Automatic parallel tiled code generation based on dependence approximation," in *Parallel Computing Technologies*, V. Malyskin, Ed. Cham: Springer International Publishing, 2021, pp. 260–275.
- [11] U. Bondhugula *et al.*, "A practical automatic polyhedral parallelizer and locality optimizer," *SIGPLAN Not.*, vol. 43, no. 6, pp. 101–113, Jun. 2008. [Online]. Available: <http://pluto-compiler.sourceforge.net>
- [12] J. Xue, *Loop Tiling for Parallelism*. Norwell, MA, USA: Kluwer Academic Publishers, 2000. ISBN 0-7923-7933-0
- [13] M. Palkowski and W. Bielecki, "NPDP benchmark suite for the evaluation of the effectiveness of automatic optimizing compilers," *Parallel Computing*, vol. 116, p. 103016, Jul. 2023. doi: 10.1016/j.parco.2023.103016. [Online]. Available: <https://doi.org/10.1016/j.parco.2023.103016>
- [14] A. Mehdi, *Par4All User Guide*, 2012. [Online]. Available: <http://www.par4all.org>
- [15] C. Dave, H. Bae, S.-J. Min, S. Lee, R. Eigenmann, and S. Midkiff, "Cetus: A source-to-source compiler infrastructure for multicores," *Computer*, vol. 42, pp. 36–42, 2009.
- [16] S. Verdoolaege, J. Carlos Juega, A. Cohen, J. Ignacio Gómez, C. Tenllado, and F. Catthoor, "Polyhedral parallel code generation for cuda," *ACM Transactions on Architecture and Code Optimization*, vol. 9, no. 4, p. 1–23, Jan. 2013. doi: 10.1145/2400682.2400713. [Online]. Available: <http://dx.doi.org/10.1145/2400682.2400713>
- [17] "Nvidia corporation, cuda programming guide 12.3," <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>, 2023, accessed on: 2023-10-22.
- [18] K. Thouti and S. R. Sathe, "Comparison of openmp & opencl parallel processing technologies," 2012. doi: 10.48550/ARXIV.1211.2038. [Online]. Available: <https://arxiv.org/abs/1211.2038>
- [19] M. Khalilov and A. Timoveev, "Performance analysis of cuda, openacc and openmp programming models on tesla v100 gpu," *Journal of Physics: Conference Series*, vol. 1740, no. 1, p. 012056, Jan. 2021. doi: 10.1088/1742-6596/1740/1/012056. [Online]. Available: <http://dx.doi.org/10.1088/1742-6596/1740/1/012056>
- [20] G. Kan, X. He, L. Ding, J. Li, K. Liang, and Y. Hong, "A heterogeneous computing accelerated sce-ua global optimization method using openmp, opencl, cuda, and openacc," *Water Science and Technology*, vol. 76, no. 7, p. 1640–1651, Jun. 2017. doi: 10.2166/wst.2017.322. [Online]. Available: <http://dx.doi.org/10.2166/wst.2017.322>
- [21] L. Chen, P.-H. Lin, T. Vanderbruggen, C. Liao, M. Emani, and B. de Supinski, *LM4HPC: Towards Effective Language Model Application in High-Performance Computing*. Springer Nature Switzerland, 2023, p. 18–33. ISBN 9783031407444. [Online]. Available: http://dx.doi.org/10.1007/978-3-031-40744-4_2
- [22] W. Godoy, P. Valero-Lara, K. Teranishi, P. Balaprakash, and J. Vetter, "Evaluation of openai codex for hpc parallel programming models kernel generation," in *Proceedings of the 52nd International Conference on Parallel Processing Workshops*, ser. ICPP-W 2023. ACM, Aug. 2023. doi: 10.1145/3605731.3605886. [Online]. Available: <http://dx.doi.org/10.1145/3605731.3605886>
- [23] G. C. Team, "Github copilot," <https://copilot.github.com/>, 2022, an AI pair programmer for GitHub, Accessed on: 2023-10-22.
- [24] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba, "Evaluating large language models trained on code," <https://arxiv.org/abs/2107.03374>, 2021, accessed on: 2023-10-22.
- [25] P. Valero-Lara, A. Huante, M. A. Lail, W. F. Godoy, K. Teranishi, P. Balaprakash, and J. S. Vetter, "Comparing llama-2 and gpt-3 llms for hpc kernels generation," 2023. [Online]. Available: <https://arxiv.org/abs/2309.07103>
- [26] "Introducing llama 2, the next generation of our open source large language model," <https://ai.meta.com/llama/>, 2023, accessed on: 2023-10-22.
- [27] D. Wonnacott, T. Jin, and A. Lake, "Automatic tiling of "mostly-tileable" loop nests," in *IMPACT 2015: 5th International Workshop on Polyhedral Compilation Techniques, At Amsterdam, The Netherlands*, 2015.
- [28] S. Verdoolaege, "Integer set library - manual," www.kotnet.org/~skimo/isl/manual.pdf, 2011, accessed on: 2024-01-11.
- [29] U. Bondhugula *et al.*, "A practical automatic polyhedral parallelizer and locality optimizer," *SIGPLAN Not.*, vol. 43, no. 6, pp. 101–113, Jun. 2008. doi: 10.1145/1379022.1375595
- [30] M. Palkowski and W. Bielecki, "Parallel tiled Nussinov RNA folding loop nest generated using both dependence graph transitive closure and loop skewing," *BMC Bioinformatics*, vol. 18, no. 1, p. 290, 2017. doi: 10.1186/s12859-017-1707-8
- [31] M. Palkowski and M. Gruzewski, "Time and energy benefits of using automatic optimization compilers for NPDP tasks," *Electronics*, vol. 12, no. 17, p. 3579, Aug. 2023. doi: 10.3390/electronics12173579. [Online]. Available: <http://dx.doi.org/10.3390/electronics12173579>

Impact of Local Geometry on Methods for Constructing Protein Conformations

W. Da Rocha,^{*} T.E. Malliavin,[†] A. Mucherino,[‡] L. Liberti^{*}

^{*}LIX, École Polytechnique, Palaiseau, France.

{ wagner.rocha, liberti } @lix.polytechnique.fr

[†]LPCT-CNRS and University of Lorraine, Vandoeuvre-lès-Nancy, France.

therese.malliavin@univ-lorraine.fr

[‡]IRISA, University of Rennes, Rennes, France.

antonio.mucherino@irisa.fr

Abstract—The prediction of protein structures is an important problem in molecular biology. In spite of the large efforts from the research community, and of the recent development of artificial intelligence tools specifically designed for this problem, a complete and definitive solution to the problem has not been found yet. This work is based on the observation that many tools for the prediction of protein conformations rely on both local and non-local geometrical information, even though the non-local information can be very hard to identify within the desired precision in some particular situations. For this reason, we explore in this work the effect of local geometry on methods capable of constructing protein conformations. This initial study has the final aim of devising new alternative methods where the predictions may be guided mainly by the local geometry of proteins.

I. INTRODUCTION

THE prediction of suitable conformations for a given protein is of fundamental importance in science, in particular in the context of drug design. Since several years the research community has been working on this topic, with the final goal of understanding how these three-dimensional conformations can be “predicted” by using some of the information that can be obtained through experimental techniques. When the predictions rely solely on the protein sequence (i.e. on the list of *amino acids* forming the main protein chain), then it is common to refer to the problem of identifying these possible conformations as the “protein folding” problem [5].

In spite of the large efforts in this scientific domain, the protein folding problem remained for several years, as long as general instances are concerned, among the practically intractable problems. Experimental techniques that are able of providing additional information about the molecules (and not only its amino acid sequence) were meanwhile developed, and methods and algorithms were thus proposed that are capable of determining protein conformations from these experimental data. One example, on which we have been working in the past 15 years, is given by the experimental technique based on Nuclear Magnetic Resonance (NMR) [2], where a Distance

Geometry Problem (DGP) [15], [17] is formulated for the determination of the protein conformations.

These methods exploit both local structural information, as well as long-range proximity measures [12]. Local structural information can for example be deduced from the study of the chemical structure of each amino acid: if two atoms are chemically bonded, then it is possible to *guess*, in a rather precise way, the distance separating the two atoms. Force fields such as AMBER [4] and PARALLHDG [6] collect a certain number of parameters which also comprise this kind of local proximity information. Naturally, the given values for such parameters are not wholly satisfied in all proteins. In fact, terms of energy functions given by such force fields are actually able to give a measure on the variations of these values in protein conformations.

It is common to talk about long-range proximity measures when we can obtain estimates on the distances between two atoms belonging to two amino acids that may be separated by several other amino acids in the protein sequence. The experiments based on NMR techniques (already mentioned above), for example, are able to give estimates on such long-range distances, and most commonly between pairs of hydrogen atoms [8]. Alternatively, methods based on multiple sequence alignments [21] can also provide long-range distance information, but they are likely to give imprecise results for particular cases of proteins [20]. A real challenge for both NMR experiments and methods based on protein sequence alignments are the so-called Intrinsically Disordered Proteins (IDPs) [19].

Nevertheless, AlphaFold (the release of the version 3 is very recent, see [1]), the very well-known Artificial Intelligence (AI) tool for protein folding, strongly relies on long-range proximity information, normally obtained from protein sequence alignments. The success of AlphaFold is therefore strongly dependent on the availability of such long-range restraints, and its actual success rate can therefore depend on the number of alignments that it is possible to exploit in order to derive the long-range proximity measures.

In this work, we intend investigating the impact of local geometry in the determination of protein folds. Our work is motivated by some previous analysis performed by some of us [9] where we have identified some particular situations in which the local geometry seems to have a larger impact on the protein folds than long-range distances (which were identified though NMR experiments in that work). Some initial investigations in line with the present work were already conducted and published in [10]. Our work extends those initial studies and uses a larger subset of protein conformations in the computational experiments.

The remainder of the paper is organized as follows. In Section II, we will describe in more detail what we intend by local geometry of protein conformations, and we will explain how to define DGP instances for protein conformations carrying specific local geometry information. In Section III, we will briefly describe the Branch-and-Prune (BP) [14], used for solving the artificially generated DGP instances. Finally, Section IV will present our preliminary computational experiments, and Section V will briefly conclude the paper.

II. LOCAL GEOMETRY OF PROTEINS

Proteins are defined by one or more sequences of amino acids. In this work, we focus our attention on proteins defined by only one amino acid chain. Amino acids are the building blocks of proteins. There are 20 different amino acids that can be involved in the protein synthesis, and they all have a common part, while another, named the *side chain* of the amino acid, makes each amino acid different from one another. The subset of atoms forming the protein which are not included in the side chains is generally referred to as the *protein backbone*.

The chemical composition for every of the 20 amino acids is a priori known, and therefore the chemical composition of the entire protein can be simply obtained from the amino acid sequence. Part of the local geometry can be as a consequence derived from a simple analysis of the atomic bonds that are present in the structure. As already mentioned, bonded atoms satisfy a relative distance (a “bond length”) that is generally considered to depend solely on the type of the two involved atoms.

Similarly, we can extend the same idea to the angle that every triplet of bonded atoms can form (say the atoms are A, B and C, where A is bonded to B, B is bonded to C, and we are interested in the angle in B formed by the segments AB and BC). In this case, we rather talk about “bond angles”, and it is again generally supposed that these angles basically depend on only the type of involved atoms. We point out, however, that in the case of bond angles, a larger variation of the angles can be observed around their average value.

The situation is a little more complex when quadruplets of consecutive atoms are defined. They allow us to define the so-called *torsion angles*. Torsion angles exhibit larger variations over the protein conformations, and they are not as regular as bond lengths and angles. However, there are some special cases where we can constrain the values of these angles. One

example is given by the torsion angle ω crossing a peptide bond (from the C_α Carbon of the amino acid i to the C_α Carbon of the amino acid $i + 1$ in the sequence), which is generally fixed and set to 178° . Another example is given by the protein secondary structures, which strongly restrict the ranges of the torsion angles for every quadruplet of atoms that we can define on the protein backbones. This is a very important result for protein conformations, which was studied for the first time in the well-known Ramachandran map [18]. In the following, we will employ the typical notations ϕ , ψ and ω for the torsion angles that we can define on the protein backbones.

In this work, we investigate the dependence of local geometry in proteins (bond lengths and angles, as well as torsion angles) on parameters other than the simple atom type, as it was instead supposed in the seminal works of Engh and Huber [6]. In particular, we will compare the three following **setups**:

1. local geometry is unique for every protein and cannot be predicted by the analysis of the protein sequence;
2. local geometry can be predicted by using the information about the secondary structures related to the protein chain, together with the atom type;
3. local geometry can be predicted by the simple analysis of the atom types (as in Engh and Huber’s works).

In order to study and compare these three setups, we will artificially generate three different sets of DGP instances, which we will solve by the BP algorithm briefly summarized in Section III. Our computational experiments will be then presented and commented in Section IV.

III. AN IMPLEMENTATION OF THE BP ALGORITHM

In this section, we will first of all introduce the DGP in formal terms, and we will briefly describe a well-known algorithm for the solution of DGP instances that can be *discretized*, and finally mention to the specific implementation of the algorithm that we will use in our computational experiments.

Let $G = (V, E, d)$ be a simple weighted undirected graph, where vertices represent the atoms of our proteins, and the existence of an edge between two atoms indicate that their relative distance is known [15]. The weight function d associates the numerical value of the distance to every edge of E . This numerical value $d(u, v)$ can be either exact (i.e. very precise), so that it can be represented by a singleton, or rather imprecise and hence represented by a real-valued interval $[d(u, v), \bar{d}(u, v)]$. Let E' be the subset of the edge set E containing only the exact distances.

Given a simple weighted undirected graph $G = (V, E, d)$, the Distance Geometry Problem (DGP) in dimension 3 asks whether a graph embedding

$$x : v \in V \longrightarrow x_v \in \mathbb{R}^3$$

exists such that

$$\forall \{u, v\} \in E, \quad \|x_u - x_v\| \in d(u, v), \quad (1)$$

where $\|\cdot\|$ represents the Euclidean norm. We say that the graph embedding x is a *realization* of the graph when it satisfies all the constraints in Eq. (1).

In the past years, some of us have been focusing on a special class of DGP instances where the search space can be discretized and reduced to a tree [13]. Let $G[\cdot]$ be the subgraph of G induced by a subset of vertices of V . In formal terms, a given DGP instance can be discretized (so that it represents an instance of the Discretizable DGP, or DDGP) when there exists a vertex ordering on V such that the following two assumptions are satisfied:

- (a) $G[\{1, 2, 3\}]$ is a clique whose edges are in E' ;
- (b) $\forall v \in \{4, \dots, |V|\}$, there exist $u_1, u_2, u_3 \in V$ such that
 - (b.1) $u_1 < v, u_2 < v, u_3 < v$;
 - (b.2) $\{\{u_1, v\}, \{u_2, v\}\} \subset E', \{u_3, v\} \in E$;
 - (b.3) $d(u_1, u_3) < d(u_1, u_2) + d(u_2, u_3)$.

When the two assumptions (a) and (b) are satisfied, we can construct a search tree where the candidate positions for every atom are collected on a common tree layer [15]. In our experiments, we consider the vertex ordering defined in [22], which makes an extensive use of repeated vertices in order to achieve a direct branching on the torsion angles ϕ , ψ and ω that we can define in the protein backbones. For strict enough values for these torsion angles, it is in fact possible to avoid branching and hence locally reduce, *a priori*, the tree width.

When the edge $\{u_3, v\}$ is not in E' (see assumption (b.2)), the distance $d(u_3, v)$ is represented by an interval. In this situation, the set of possible positions for the atom v is actually continuous, but in some particular conditions (which are satisfied by the instances we use in this work) we can consider to take sample distance values from the original intervals, and to have a dedicated branch in the tree for every extracted sample distance [13]. This methodology introduces an additional factor (given by the number of samples taken from every interval distance) in the combinatorics, but it has the advance to make us deal with more complex instances by using an approach that was initially designed to work in simpler conditions (i.e. with distances that are not affected by uncertainty).

The Branch-and-Prune (BP) algorithm performs a systematic exploration of this search tree. It uses the additional distances, which are not necessary for the construction of the tree, to verify the “feasibility” of the generated atomic positions [14]. This is the so-called *pruning phase* of the algorithm, which is actually very important in BP, because it allows the algorithm to focus the search over the tree branches that contain no infeasibilities.

For more information about the BP algorithm and its previous uses in the context of structural biology, the reader is mainly referred to [7], [16]. In our computational experiments, we will use the implementation of the BP algorithm available on the following GitHub repository:

<https://github.com/tmalliavin/ibp-ng-fullchain>

IV. COMPUTATIONAL EXPERIMENTS

In our computational experiments, we have selected a subset of protein conformations from the Protein Data Bank (PDB) [3]. The conformations have been selected in order to satisfy the following properties:

- the conformations are obtained through techniques that are based on X-ray crystallography, with resolution of at least 1.6 Å and crystallographic R factor larger than 0.25;
- the protein sequences (only one chain) are not longer than 100 amino acids;
- the similarity between the amino acid sequences of any pair of proteins is smaller than 20%;
- the molecules do not contain *cis* peptide bonds;
- at least two secondary structure elements (α -helix or β -strand) are present in the protein.

Our subset finally contains 308 protein conformations, and we consider the three main setups listed in Section II for the generation of our artificial instances. When we use **setup 1**, we suppose that the local geometry is unique for every protein, and hence we extract the information from the original PDB conformations. When we use **setup 2**, the distances and angles that we use to define our instances are averaged over the Hollingsworth’s regions [11], which provide a finer-grained partitioning of initial Ramachandran regions. Finally, we use the values proposed by Engh and Huber’s works [6] under the hypothesis that they can only depend on the atom types (our **setup 3**). Notice that these setups can be mixed so that a different one can be considered for a different kind of local information. Details for each set of performed experiments are given in the caption of Fig. 1.

The protein conformation have been reconstructed using “one-shot” BP runs. The branching phase is performed with a discretization factor allowing to have a variation on the torsion angles ϕ and ψ of magnitude about 5° . The ω values are instead used in the pruning phase. A scaling factor of 0.8 is applied to van der Waals radii in order to introduce lower bounds on unknown distances (atoms cannot be closer than a certain threshold when they are not chemically bonded), and the error tolerance ϵ is set to 0.1. The runs are stopped as soon as the first solution is constructed.

Fig. 1 displays the distributions of the root-mean-square deviation (RMSD, Å) between the atomic coordinates related to the solutions found by the BP algorithm and the original PDB conformations. As expected, the best results in terms of RMSD are obtained when the local geometry is extracted from the original PDB conformations (Fig. 1(a)). Interestingly, when we consider the bond lengths from Engh and Huber’s works, a similar distribution is obtained (data not shown), implying therefore that the bond lengths have actually little impact on the reconstruction process for the conformations. The influence of the bond angles appears instead to be more important.

When the torsion angles ω are imposed to 178° degrees, we can notice a large increase of the RMSD values, reaching values of 10 or even 12 Å (see Fig. 1(b)). This result shows

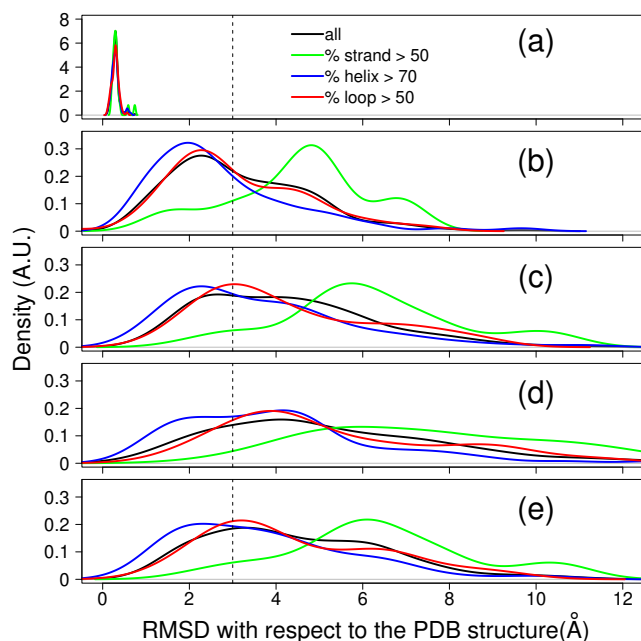


Fig. 1. Distribution of the root-mean-square deviation (\AA) between the original PDB conformations and those reconstructed using the BP algorithm. From up to down: (a) setup 3 for bond lengths, setup 1 for the remaining local geometry; (b) as previous one but with ω angles fixed to 178° ; (c) setup 2 for the entire local information; (d) as previous one but with ω angles set to 178° ; (e) setup 3 for the entire local geometry. The dashed vertical line is placed at 3 \AA .

the importance of the local variability in the peptide bond geometry.

We can remark moreover that some RMSD distributions strongly depend on the secondary structures that are actually contained in the conformations. Proteins with a large percentage of α -helices (blue lines) or mostly containing loops (red lines) are shifted towards smaller RMSD values. By contrast, the structures containing mostly β -strands (green lines) exhibit RMSD values larger than 3 \AA . The calculations of conformations based on Ramachandran regions with ω values extracted from the PDB conformations (see Fig. 1(c)) display slightly better RMSD values than the ones using the fixed value of 178° (see Fig. 1(d)).

Finally, when it is supposed that local geometry only depends on the atom types (see Fig. 1(e)), then the observed RMSD values are even larger. These results thus underline the negative impact of uniform local geometry on the reconstructed conformations.

V. CONCLUSIONS

We have presented a study on the local geometrical information of protein conformations. Even though we are aware that the local and global geometry are likely to be highly entangled in proteins, this work consisted in investigating how much the local information can have an impact on the protein folds. This study, together with others that we plan to perform in the

near future, can potentially help us in developing new methods and algorithms for the construction of protein conformations which mainly (or even solely) use information about the local geometry of proteins.

ACKNOWLEDGMENTS

This work was partially supported by the CNRS (ITINERANCE and IRP projects), Lorraine University, IRISA, ANR PRCI multiBioStruct (ANR-19-CE45-0019). High Performance Computing resources were provided by the EXPLOR center at Lorraine University (2022CPMXX2687).

REFERENCES

- [1] J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A.J. Ballard, J. Bambrick, S.W. Bodenstern, D.A. Evans, C.C. Hung, M. O'Neill, D. Reiman, K. Tunyasuvunakool, Z. Wu, A. Zengulytė, E. Arvaniti, C. Beattie, O. Bertolli, A. Bridgland, A. Cherepanov, M. Congreve, A.I. Cowen-Rivers, A. Cowie, M. Figurnov, F.B. Fuchs, H. Gladman, R. Jain, Y.A. Khan, C.M.R. Low, K. Perlin, A. Potapenko, P. Savy, S. Singh, A. Stecula, A. Thillaisundaram, C. Tong, S. Yakneen, E.D. Zhong, M. Zielinski, A. Židek, V. Bapst, P. Kohli, M. Jaderberg, D. Hassabis, J.M. Jumper, *Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3*, to appear in Nature, accelerated preview on Nature.com published on May 8, 2024.
- [2] F.C.L. Almeida, A.H. Moraes, F. Gomes-Neto, *An Overview on Protein Structure Determination by NMR, Historical and Future Perspectives of the Use of Distance Geometry Methods*. In: [17], 377–412, 2013.
- [3] H. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. Bhat, H. Weissig, I. Shindyalov, P. Bourne, *The Protein Data Bank*, *Nucleic Acids Research* **28**, 235–242, 2000.
- [4] D.A. Case, R.M. Betz, D.S. Cerutti, T.E. Cheatham III, T.A. Darden, R.E. Duke, T.J. Giese, H. Gohlke, A.W. Goetz, N. Homeyer, S. Izadi, P. Janowski, J. Kaus, A. Kovalenko, T.S. Lee, S. LeGrand, P. Li, C. Lin, T. Luchko, R. Luo, B. Madej, D. Mermelstein, K.M. Merz, G. Monard, H. Nguyen, H.T. Nguyen, I. Omelyan, A. Onufriev, D.R. Roe, A. Roitberg, C. Sagui, C.L. Simmerling, W.M. Botello-Smith, J. Swails, R.C. Walker, J. Wang, R.M. Wolf, X. Wu, L. Xiao, P.A. Kollman, *AMBER 2016*, University of California, San Francisco, 2016.
- [5] K.A. Dill, S. Banu Ozkan, M. Scott Shell, T.R. Weikl, *The Protein Folding Problem*, *Annual Review of Biophysics* **37**, 289–316, 2008.
- [6] R. Engh, R. Huber, *Accurate Bond and Angle Parameters for X-ray Protein Structure Refinement*, *Acta Crystallographica A* **47**, 392–400, 1991.
- [7] D. Förster, J. Idier, L. Liberti, A. Mucherino, J.-H. Lin, T.E. Malliavin, *Low-Resolution Description of the Conformational Space for Intrinsically Disordered Proteins*, *Scientific Reports* **12**, 19057, 16 pages, 2022.
- [8] P. Güntert, L. Buchner, *Combined Automated NOE Assignment and Structure Calculation with CYANA*, *Journal of Biomolecular NMR* **62**, 453–471, 2015.
- [9] S.B. Hengeveld, T. Malliavin, J.H. Lin, L. Liberti, A. Mucherino, *A Study on the Impact of the Distance Types Involved in Protein Structure Determination by NMR*, IEEE Conference Proceedings, Computational Structural Bioinformatics Workshop (CSBW21), International Conference on Bioinformatics & Biomedicine (BIB21), online event, 9 pages, 2021.
- [10] S.B. Hengeveld, M. Merabti, F. Pascale, T.E. Malliavin, *A Study on the Covalent Geometry of Proteins and Its Impact on Distance Geometry*, Lecture Notes in Computer Science **14072** (part 2), F. Nielsen, F. Barbaresco (Eds.), Proceedings of Geometric Science of Information (GSI23), Saint Malo, France, 520–530, 2023.
- [11] S.A. Hollingsworth, M.C. Lewis, D.S. Berkholz, W.K. Wong, P.A. Karplus, *(phi,psi) Motifs: a Purely Conformation-based Fine-Grained Enumeration of Protein Parts at the Two-Residue Level*, *Journal of Molecular Biology* **416**(1), 78–93, 2012.
- [12] B. Kuhlman and P. Bradley, *Advances in protein structure prediction and design*, *Nature Reviews Molecular Cell Biology* **20**, 681–697, 2019.

- [13] C. Lavor, L. Liberti, A. Mucherino, *The interval Branch-and-Prune Algorithm for the Discretizable Molecular Distance Geometry Problem with Inexact Distances*, *Journal of Global Optimization* **56**(3), 855–871, 2013.
- [14] L. Liberti, C. Lavor, N. Maculan, *A Branch-and-Prune Algorithm for the Molecular Distance Geometry Problem*, *International Transactions in Operational Research* **15**, 1–17, 2008.
- [15] L. Liberti, C. Lavor, N. Maculan, A. Mucherino, *Euclidean Distance Geometry and Applications*, *SIAM Review* **56**(1), 3–69, 2014.
- [16] T.E. Malliavin, *Tandem Domain Structure Determination based on a Systematic Enumeration of Conformations*, *Scientific Reports* **11**, 16925, 2021.
- [17] A. Mucherino, C. Lavor, L. Liberti, N. Maculan (Eds.), *Distance Geometry: Theory, Methods and Applications*, 410 pages, Springer, 2013.
- [18] G.N.T. Ramachandran, V. Sasisekharan, *Conformation of Polypeptides and Proteins*, *Advances in Protein Chemistry* **23**, 283–437, 1968.
- [19] P. Tompa, *Intrinsically Disordered Proteins: a 10-Year Recap*, *Trends in Biochemical Sciences* **37**(12), 509–516, 2012.
- [20] T. Warnow, *Revisiting Evaluation of Multiple Sequence Alignment Methods*, *Methods in Molecular Biology* **2231**, 299–317, 2021.
- [21] M. Weigt, R.A. White, H. Szurmant, J.A. Hoch, T. Hwa, *Identification of Direct Residue Contacts in Protein-Protein Interaction by Message Passing*. *PNAS* **106**, 67–72, 2009.
- [22] B. Worley, F. Delhommel, F. Cordier, T.E. Malliavin, B. Bardiaux, N. Wolff, M. Nilges, C. Lavor, L. Liberti, *Tuning Interval Branch-and-Prune for Protein Structure Determination*, *Journal of Global Optimization* **72**, 109–127, 2018.

The comparison of pixel-based image analysis for detection of weeds in winter wheat from UAV imagery

Vojtěch Slezák
0009-0007-8961-7793
Mendel University in Brno
Department of Agrosystems
and Bioclimatology
Zemědělská 1, 613 00
Czech Republic
Email: xslezak2@mendelu.cz

Kateřina Kuchaříková,
0009-0002-7285-5553
Tomáš Kaplánek
0009-0007-7543-1090

Vojtěch Lukas,
0000-0001-8051-3305
Jan Křen
0000-0002-5229-398X

Abstract—Creating weed maps directly by growers is becoming increasingly common. In this study, an unmanned aerial vehicle (UAV) imaged a field infested by field thistle (*Cirsium arvense*). This paper compares four detection methods that can be used concerning agricultural practice. Two algorithms are supervised classification methods - Maximum Likelihood (ML) and Supported Vector Machine (SVM). The Pix4Dfields (Magic Tool) classification algorithm and the thresholding method are other methods used. The Kappa coefficient and the overall accuracy determined the accuracy of the individual algorithms. The highest accuracy was achieved by the thresholding method, and the lowest by the Pix4Dfields algorithm. Among the supervised classification methods, SVM achieved higher accuracy than the ML algorithm. In terms of using the methods in practice, the thresholding method proved more effective than supervised classification methods.

Index Terms— Precision agriculture, SSWM, Pix4D, Remote sensing

I. INTRODUCTION

OPTIMISING the use of herbicides is a primary goal and is crucial for maintaining the competitiveness of farms and decreasing the consumption of agrochemicals in crop management. Effective site-specific weed management (SSWM) requires knowledge of the spatial variability of the weed plants in the field. UAVs capable of capturing images with high spatial and spectral resolution have proven effective for this purpose [1], [2]. RTK modules, which accurately locate both the UAV and the photos taken, are now a standard feature. Some UAVs have the advantage of carrying multiple interchangeable or integrated sensors simultaneously. Sensors that capture in the visible spectrum (RGB) and multispectral cameras that acquire data from red-edge (RE) or near-infrared (NIR) bands are commonly used for weed detection [3], [4].

Confidence in accurate weed management is bolstered by modern application technology of sprayers that allows section-by-section or individual nozzle control based on the pre-

scription map. This map typically consists of polygons outlining the application area [5], where nozzles are turned on. The accuracy of the application is ensured by the RTK guidance systems mounted on the machinery [6].

The technique for detecting weed plants from acquired images depends on the sensor type, the computing technology's performance, and the spraying technique. The modern sprayers with individual nozzle control require higher accuracy of prescription maps than older sprayers with section application swath control [5]. Besides the section control, the sprayer terminal's computing technology is crucial to processing the large data sets of the detailed prescription map. Weed detection in realistic conditions uses images combined into a single orthorectified mosaic (orthomosaic) representing the entire area of interest. Commercial software like Pix4D fields, Agisoft Metashape, Drone Deploy, or open-source options like OpenDroneMap are commonly used for this purpose [7], [8].

In addition to site-specific applications using a prescription map derived from UAV imagery, real-time detection and application techniques are also used. Techniques for detecting green vegetation on bare soil based on spectral features have existed since the last century [9]. Currently, sensor systems on application technology can detect weeds in broad-row crops like corn and soybeans. These systems save the cost of field surveys by UAV imaging but are usually more expensive than section control sprayers and can only be used in specific cases [10].

Application drones capable of applying solid and liquid products are increasingly becoming an alternative to traditional boom sprayers. The drone system automatically suggests a flight path and the optimal flight level for the application. This allows pesticide application in conditions where conventional machinery sprayers cannot enter the field, such as in unsuitable soil conditions, to avoid soil destruction and compaction [11]. Application drones are considerably cheaper than conventional sprayers, making them an attractive option for some growers. However, their weakest link is

the batteries, which allow flight times of around 10 minutes, with excessive heating during charging slowing down the process [10], [12].

Identification of weed infestation based on the UAV data

There are two main categories of weed detection from UAV image data: object-based image analysis (OBIA) and pixel-based image analysis (PBIA). Supervised classification on high spatial resolution images taken by UAVs has promising results [13]. Supervised PBIA classifiers leverage prior knowledge to identify spectral similarities in raster data, assigning each pixel to the most appropriate class. In PBIA for supervised weed classification, each pixel may contain a mix of soil, plant leaves, residue, and shadow. This mixture can introduce variability in the target's spectral reflectance, limiting classification accuracy. In high-resolution PBIA, the heterogeneity of spectral values within a single class further restricts its effectiveness [14]. Compared to these two approaches, OBIA usually achieves better results than PBIA. Each classifier has its limitations, and the choice of the classifier depends on many factors, such as the data's spectral and spatial resolution, classification accuracy, algorithm performance, and computational resources [13; 15].

The study aimed to validate weed detection techniques on images taken by UAVs and concerning agricultural practices. Four classifiers were chosen for verification: two basic classifiers based on machine learning: Maximum Likelihood (ML), which works best when class samples are normally distributed, and Supported Vector Machine (SVM), which is commonly used in the research community and can handle standard images as well as segmented images, with less susceptible to noise, correlated bands [16]. The other two classifiers used in this study are a Pix4Dfields (Magic Tool) classification tool and a pixel extraction procedure based on vegetation index threshold setting. All classifiers were chosen for their versatility and simple hyperparameter adjustment, which are important elements in agronomic practice.

II. MATERIAL AND METHODS

A. Study area

The study was realized in 2023 in the form of field trial with the area 15.85 located at Rataje site (Kromeriz, Czech Republic; 49.254° N, 17.332° E). The main investigation was focused on the detection of occurrence of weed "field thistle" (*Cirsium arvense*) in winter wheat in the early stage of crop growth (BBCH 10-13).




Fig. 1 Position of the field of study in the Czech Republic.

B. Data acquisition

The UAV imagery was carried out on October 12, 2023, by the quadcopter DJI Matrice 300 RTK. The drone was equipped with two sensors: RGB sensor DJI Zenmuse P1 and multispectral sensor MicaSense RedEdge-P. The flying altitude was 120 m above the ground, with GSD 13 mm (RGB camera) and 40 mm per pixel (RedEdge-P). Flying speed was set to 8.2 m/s. Both sensors were positioned in a nadir view of the canopy. The image overlap ratio was set at 70 % side and 80 % frontal. The RGB sensor exposure time was set at 1/2000 s, and the ISO and aperture were set to auto with timed interval shot. On multispectral sensors, images were taken every 1.5 s. After the flight mission was completed, a calibration reflection panel was photographed. The final orthomosaic from UAV imagery was processed using Pix4Dfields software.




Fig. 2 Stitched images and a sample of resolution.

C. Classification approach

Classification algorithms based on ML and SVM were triggered in the ESRI ArcGIS Pro software. The classification was performed on RGB imagery. Due to the crop's early phenological stage and low spatial resolution of orthomosaic the detection of winter wheat plants was limited. Thus, only two classes were identified - bare soil (mixed with winter wheat plants) and weeds (as full vegetation cover). The OBIA

method was not used in this case, and the classification was made using PBIA techniques [17]. The PBIA classifications were performed on the original image. Algorithms ML and SVM were trained using samples created by an expert based on ground truth identification. The training dataset included 50 samples for each category.

The Pix4Dfields mapping software includes a "Magic tool" classification tool. It is a simple and user-friendly multipurpose classification tool based on machine learning, which can detect weed patches or plant damage. The algorithm uses a grid over the field and the user labels grids that should be treated and untreated. The recommended number of labelled grids per class is 20; the minimum is three. In our case, the size of the grid over the field was set to 1 m and rotated in the driving direction. Subsequently, 40 grids (20 per class) were selected as the training samples.




Fig. 3 Example of grid labelling using Pix4Dfields classification tool.

The significant difference in vegetation phases between the main crop and the detected weeds indicates a difference in spectral characteristics, which offers the possibility of using the thresholding method in the image. Thresholding is a technique to segment images by creating binary images based on threshold settings. The input raster was the vegetation index NDVI (1) of the studied field.

$$NDVI = \frac{(\rho_{NIR} - \rho_{Red})}{(\rho_{NIR} + \rho_{Red})} \quad (1)$$

The threshold value was set based on the histogram of the distribution of data values to value 0.385 as the cut-off between crops with soil and weeds. Reclassifying produced a binary image and left the pixels that matched the weeds. Pixels of the raster were converted to vector points, and a weed coverage map was created using a buffer tool with a buffer distance of 0.1 m. The final step was dissolving buffer zones to create a polygon map, which was unified by joining the overlapping polygons.




Fig. 4 Example of thresholding NDVI.

The Accuracy Assessment Tool was used with ArcGIS Pro software to estimate the accuracy of classification algorithms. In total, 500 randomly stratified samples. These points contain classified values of all algorithms and the ground truth value based on visual verification. Based on that, a Confusion Matrix was calculated to determine the accuracy of each algorithm. Algorithm accuracy was expressed based on the overall accuracy (OA) and the Kappa coefficient [18]. OA quantifies the level of agreement between two classes (weed and soil), and it is calculated from a number of True positive (TP) samples, which represents classification that matches with the truth and from True negative (TN) samples, which were misclassified. Using (1), the OA was calculated. The percentage ranges from 0-100 %; a higher number indicates a more accurate classification [19].

$$OA = \frac{True\ Positive + True\ Negative}{Number\ of\ samples} \times 100 \quad (2)$$

Meanwhile, the Kappa coefficient represents the level of agreement between two classes corrected by chance. Kappa takes into account the number of samples that are assigned to each class. If the validation points are predominantly represented in one class, the OA will be higher regardless of the number of elements in the other. The level of the Kappa coefficient ranges from 0-1 and provides information if the classifier is better or worse than by random chance. A higher number indicates a more accurate classification. The Kappa coefficient is calculated using (3) from OA and Chance Agreement, which is calculated as the sum of the product of row and column totals for each class [19], [20].

$$Kappa\ Coefficient = \frac{OA + chance\ agreement}{1 - chance\ agreement} \quad (3)$$

III. RESULT AND DISCUSSION

The results of Accuracy Assessment Tool are presented in Table I. The accuracy of the Maximum Likelihood algorithm

was 97.8 %, and the Kappa coefficient was 0.633 with an area of 3,643 m² as a weed detected. In contrast, the SVM algorithm performed better with an accuracy of 98.2 %, and the Kappa coefficient was 0.792, an area of weed coverage map increase of 5,233 m². When comparing those two algorithms, SVM produced more accurate results than ML. In terms of OA, both models performed well. However, the higher Kappa coefficient by SVM indicates a more accurate one. This confirms that SVM performs better in small training sample sizes than other models [19].

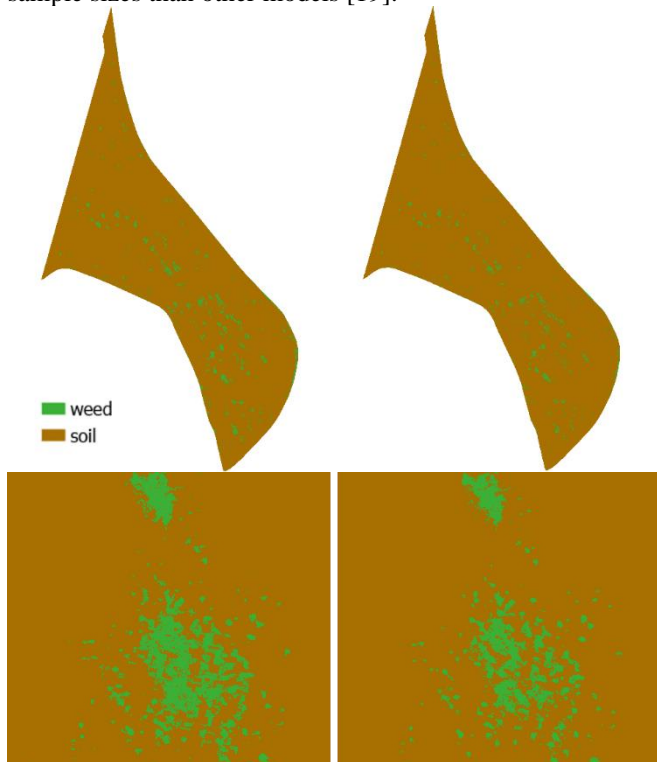


Fig. 5 Classified image by SVM algorithm on the left and ML on the right with the class of weed (green) and mixed bare soil and crop (winter wheat in early stage, brown color)

The detection algorithm by Pix4Dfields software achieved the lowest accuracy; the evaluated accuracy was 96.2 %, but the Kappa index was 0.599. The weed coverage was 5.01 %, representing an area of 7,942 m². However, we must remember that the algorithm runs in a square grid (1 m), leading to a higher misleading value. Even though the algorithm achieved the lowest accuracy, the software is applicable in practice and often used by farmers. If a detected weed is even partially in the square, it is marked as a detected weed. This leads to a larger detected area and a higher probability that the algorithm's accuracy based on the point validation data will be lower. After automatic evaluation of the detection, cleaning up the map square by square is possible, which can lead to a good result. Due to the inability to save the training dataset, it is time-consuming to label the grids for each field and then manually clean up the map.

The thresholding method achieved the highest accuracy, with an overall accuracy of 98.6% and a Kappa index of 0.836; the detected area was 8,817 m². The accuracy of the vegetation index thresholding method depends on the precise

determination of the threshold value, which is determined subjectively based on the distribution of histogram values [21]. This method was the most accurate of the four used in this comparison. The detected area was the highest, and this size can be affected by setting the distance in the buffer tool. Using the buffer tool with dissolve result will simplify the map; the larger the buffer distance, the fewer the polygons and the smaller the map size. However, with a larger buffer distance, the savings from herbicide application are reduced.

TABLE I.
VALUES OF OVERALL ACCURACY, KAPPA COEFFICIENT AND WEED COVERAGE IN THE INVESTIGATED METHODS

Method	Overall accuracy	Kappa coefficient	Weed coverage
ML	97.8 %	0.633	2.29 %
SVM	98.2 %	0.792	3.30 %
Pix4Dfields	96.2 %	0.599	5.01 %
Thresholding	98.6 %	0.836	5.56 %

Creating a weed coverage map for multiple parcels using Pix4Dfields software is time-consuming. Still, the classification tool is intuitive, and even if it does not achieve such accuracy, there is an option to clean up the map. The thresholding method requires basic knowledge of geoinformation software (GIS). It places higher demands on the user's knowledge, but it presents a fast and accurate way of creating weed maps that can be automated to some extent. From the two machine learning-based methods, the SVM is more accurate than ML, but both have a lower detected area than the other methods tested in this paper. However, they require more advanced knowledge of GIS, model parameter settings, and sample collection training. From this perspective, thresholding methods and the Pix4Dfields classifier are more suitable for farmers. However, it also depends on other factors, such as the area of the detected plots, the knowledge of the detection procedures and the software used, and the time possibilities of the grower.

IV. CONCLUSION

The results indicate that the thresholding method achieved the highest accuracy based on the Kappa coefficient and overall accuracy, while Pix4Dfields had the lowest. Among the supervised methods, SVM outperformed ML.

The identification of weed infestation by UAV imagery has shown that only a small part of the field area is covered by weeds (up to 5.56 %). Thus, site-specific spraying can significantly reduce the amount of herbicides. From the four verified weed detection algorithms, the thresholding achieved the highest accuracy (98.6%, Kappa index 0.836). However, the detected area of the weed occurrence was the highest (5.56 %).

Reference [22] confirms that saving herbicides could be more than 90 % in specific scenarios. This offers the opportunity to target herbicides in an environmentally friendly way, with less impact on crops and lower costs.

Further research will be necessary to optimize the parameter settings of the algorithms and to verify their effectiveness under different conditions. This would allow more accurate identification and localization of weeds, thereby increasing the efficiency of herbicide application and minimizing their negative impact on the environment. Future work will also incorporate other different algorithms to further increase the accuracy and reliability of weed detection.

ACKNOWLEDGMENT

The study was supported by the Internal Grant Agency of the Faculty of AgriSciences at Mendel University in Brno as the research project IGA24-AF-IP-043.

REFERENCES

- [1] N. Ubben, M. Pukrop, and T. Jarmer, "Spatial Resolution as a Factor for Efficient UAV-Based Weed Mapping—A Soybean Field Case Study", *Remote Sensing*, vol. 16, no. 10, 2024.
- [2] J. Su, D. Yi, M. Coombes, C. Liu, X. Zhai, K. McDonald-Maier, and W. -H. Chen, "Spectral analysis and mapping of blackgrass weed by leveraging machine learning and UAV multispectral imagery", *Computers and Electronics in Agriculture*, vol. 192, 2022.
- [3] T. B. Shahi, S. Dahal, C. Sitaula, A. Neupane, and W. Guo, "Deep Learning-Based Weed Detection Using UAV Images: A Comparative Study", *Drones*, vol. 7, no. 10, 2023.
- [4] G. Castellano, P. De Marinis, and G. Vessio, "Weed mapping in multispectral drone imagery using lightweight vision transformers", *Neurocomputing*, vol. 562, 2023.
- [5] V. Vijayakumar, Y. Ampatzidis, J. K. Schueller, and T. Burks, "Smart spraying technologies for precision weed management: A review", *Smart Agricultural Technology*, vol. 6, 2023.
- [6] S. Meesaragandla, M. P. Jagtap, N. Khatri, H. Madan, and A. A. Vadduri, "Herbicide spraying and weed identification using drone technology in modern farms: A comprehensive review", *Results in Engineering*, vol. 21, 2024.
- [7] C. de Villiers, C. Munghezulu, Z. Mashaba-Munghemezulu, G. J. Chirima, and S. G. Tesfamichael, "Weed Detection in Rainfed Maize Crops Using UAV and PlanetScope Imagery", *Sustainability*, vol. 15, no. 18, 2023.
- [8] S. Villette, T. Maillot, J. -P. Guillemain, and J. -P. Douzals, "Assessment of nozzle control strategies in weed spot spraying to reduce herbicide use and avoid under- or over-application", *Biosystems Engineering*, vol. 219, pp. 68-84, 2022.
- [9] R. Raja, T. T. Nguyen, D. C. Slaughter, and S. A. Fennimore, "Real-time weed-crop classification and localisation technique for robotic weed control in lettuce", *Biosystems Engineering*, vol. 192, pp. 257-274, 2020.
- [10] M. Spaeth, M. Sökefeld, P. Schwaderer, M. E. Gauer, D. J. Sturm, C. C. Delatrée, and R. Gerhards, "Smart sprayer a technology for site-specific herbicide application", *Crop Protection*, vol. 177, 2024.
- [11] S. Meesaragandla, M. P. Jagtap, N. Khatri, H. Madan, and A. A. Vadduri, "Herbicide spraying and weed identification using drone technology in modern farms: A comprehensive review", *Results in Engineering*, vol. 21, 2024.
- [12] L. Mariga, I. Silva Tiburcio, C. A. Martins, A. N. Almeida Prado, and C. Nascimento, "Measuring battery discharge characteristics for accurate UAV endurance estimation", *The Aeronautical Journal*, vol. 124, no. 1277, pp. 1099-1113, 2020.
- [13] A. Shirzadifar, S. Bajwa, J. Nowatzki, and A. Bazrafkan, "Field identification of weed species and glyphosate-resistant weeds using high resolution imagery in early growing season", *Biosystems Engineering*, vol. 200, pp. 200-214, 2020.
- [14] T. Blaschke, G. J. Hay, M. Kelly, S. Lang, P. Hofmann, E. Addink, R. Queiroz Feitosa, F. van der Meer, H. van der Werff, F. van Coillie, and D. Tiede, "Geographic Object-Based Image Analysis – Towards a new paradigm", *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 87, pp. 180-191, 2014.
- [15] H. Huang, Y. Lan, A. Yang, Y. Zhang, S. Wen, and J. Deng, "Deep learning versus Object-based Image Analysis (OBIA) in weed mapping of UAV imagery", *International Journal of Remote Sensing*, vol. 41, no. 9, pp. 3446-3479, May 2020.
- [16] J. -L. TANG, D. -J. HE, X. JING, and F. David, "Maize seedling/weed multiclass detection in visible/near infrared image based on SVM", *JOURNAL OF INFRARED AND MILLIMETER WAVES*, vol. 30, no. 2, pp. 97-103, Mar. 2011.
- [17] N. Ubben, M. Pukrop, and T. Jarmer, "Spatial Resolution as a Factor for Efficient UAV-Based Weed Mapping—A Soybean Field Case Study", *Remote Sensing*, vol. 16, no. 10, 2024.
- [18] N. Islam, M. M. Rashid, S. Wibowo, C. -Y. Xu, A. Morshed, S. A. Wasimi, S. Moore, and S. M. Rahman, "Early Weed Detection Using Image Processing and Machine Learning Techniques in an Australian Chilli Farm", *Agriculture*, vol. 11, no. 5, 2021.
- [19] G. Rozenberg, R. Kent, and L. Blank, "Consumer-grade UAV utilized for detecting and analyzing late-season weed spatial distribution patterns in commercial onion fields", *Precision Agriculture*, vol. 22, no. 4, pp. 1317-1332, 2021.
- [20] A. P. Nicolau, K. Dyson, D. Saah, and N. Clinton, "Accuracy Assessment: Quantifying Classification Quality", *Cloud-Based Remote Sensing with Google Earth Engine*, pp. 135-145, Oct. 2024.
- [21] Z. Wu, Y. Chen, B. Zhao, X. Kang, and Y. Ding, "Review of Weed Detection Methods Based on Computer Vision", *Sensors*, vol. 21, no. 11, 2021.
- [22] J. Elbl, V. Lukas, J. Mezera, I. Hunady, and A. Kintl, "Using Self-Propelled Sprayers For The Targeted Application Of Herbicides", pp. 307-314, Oct. 2023.

An Improved Genetic Algorithm for Set Cover using Rosenthal Potential

Dena Tayebi

University College Dublin
Ireland

Email: dena.tayebi@ucdconnect.ie
ORCID: 0000-0001-6447-7930

Saurabh Ray

New York University
Abu Dhabi

Email: saurabh.ray@nyu.edu
ORCID: 0009-0005-6708-125X

Deepak Ajwani

University College Dublin
Ireland

Email: deepak.ajwani@ucd.ie
ORCID: 0000-0001-7269-4150

Abstract—A major issue with heuristics for set-cover problem is that they tend to get stuck in a local optimum typically because a large local move is necessary to find a better solution. A recent theoretical result shows that replacing the objective function by a proxy (which happens to be Rosenthal potential function) allows escaping such local optima even with small local moves albeit at the cost of an approximation factor. The Rosenthal potential function thus has the effect of *smoothing* the optimization landscape appropriately so that local search works. In this paper, we use this theoretical insight to design a simple but robust genetic algorithm for weighted set cover. We modify the fitness function as well as the crossover operator of the genetic algorithm to leverage the Rosenthal potential function. We show empirically this greatly improves the quality of the solutions obtained especially in examples where large local moves are required.

Our results are better than existing state of the art genetic algorithms and also comparable in performance with the recent local search algorithm NuSC (carefully engineered for set cover) on benchmark instances. Our algorithm, however, performs better than NuSC on simple synthetic instances where starting from an initial solution, large local moves are necessary to find a solution that is close to optimal. For such instances, our algorithm is able to find near optimal solutions whereas NuSC either takes a very long time or returns a much worse solution.

I. INTRODUCTION

THE SET cover problem (SCP) is one of the most fundamental and well studied problem in theoretical computer science. An instance of the set cover problem consists of a ground set X and a set $\mathcal{S} = \{S_1, \dots, S_m\}$ of subsets of X . The goal is to pick the smallest subset $Y \subseteq \mathcal{S}$ so that the union of the sets in Y is equal to X . In the weighted version of the problem, each set has an associated non-negative weight and the goal is to minimize the total weight of the sets in Y .

A simple greedy algorithm [1] is known to give an $O(\log n)$ -approximation (even for the weighted variant) and under standard complexity theoretic assumptions this is asymptotically the best achievable in polynomial time [2]. However, instances of the set cover problems that arise from practical applications are often not worst case instances and one can hope to do better. For

instance, significantly better algorithms are known for geometric instances of the set cover problem [3].

A major issue with heuristics for the set-cover problem is that they tend to get stuck in a local optimum. Even for the unweighted setting, it is not difficult to construct examples in which there are local optima for which an arbitrarily large change is necessary to improve the solution. A recent paper [4] finds a way to make local search as good as the greedy algorithm for the general set cover problem (which, as we mentioned before, is the best possible in polynomial time). The main idea is to replace the objective function by a proxy called the Rosenthal potential [5]. With just this change (modulo a few technical details), a local search algorithm which only adds or removes one set from the current solution in any step yields an $O(\log n)$ approximation! While this result in itself is only of theoretical interest, we believe that the idea of changing the objective function is powerful and is likely to have a practical impact since a large number of heuristics for optimization problems are based on local search. In this paper, we present a simple approach to incorporate this idea into the Genetic Algorithm metaheuristic, which yields promising empirical results. In addition to changing the objective function, we need one more crucial ingredient: a *crossover* operator (called the *minimalization operator*) that also utilizes Rosenthal potential function. We also use an idea similar to *simulated annealing* and slowly fade away the effect of the Rosenthal potential so that in the end we are left with the original objective function for the (weighted) set cover problem.

II. RELATED WORK

Given the fundamental importance of set cover problem, a large number of algorithmic techniques have been developed for this problem. These include exact algorithms (e.g., [6]) and approximation algorithms (e.g. [7]). In addition, a number of heuristics and metaheuristics have been developed for this problem which is also the focus of this paper. A significant research focus in this area has been on unweighted set cover

algorithms that deal with instances where all sets have a uniform weight (also called unicost SCP). These include a greedy randomized adaptive search procedure (GRASP) [8], element-state configuration checking to cut down search spaces [9], a local search algorithm based on "electromagnetism" theory [10] and an adaptive row weighting algorithm [11]. In contrast, the weighted set cover heuristics often rely on MIP or MaxSAT formulations (e.g., [12]), metaheuristics (such as simulated annealing [13], genetic algorithm [14], bee colony [15]), local search (e.g., [12], [16]) and greedy heuristics (e.g. [17]).

Very recently, a local search heuristic NuSC [16] was proposed that outperforms the other heuristics on benchmark instances. This heuristic combines the strengths of algorithmic preprocessing (to reduce the search space), a greedy algorithm (for generating an initial solution), Tabu Search (to remove some local moves for consideration for a limited time) and local search moves based on a carefully designed scoring functions to add and remove subsets. We use NuSC as the state-of-the-art baseline to compare our genetic algorithm and show that while it performs very well on the benchmark instances because of its careful engineering, it is easy to generate instances where it fails to provide good solutions in a reasonable time.

Genetic algorithms is an extremely popular metaheuristic framework that is widely used in the design of optimization heuristics. We refer the reader to a recent survey [18] on the past, present and future of genetic algorithms. For set cover problem, Beasley and Chu [14] gave a genetic algorithm that has been widely used. In their genetic algorithm, they used a fitness-based fusion crossover operator, a variable mutation rate and a heuristic feasibility operator tailored specifically for the set covering problem. We use this work as one of the baselines for comparison.

III. ROSENTHAL POTENTIAL FUNCTION AND SET COVER

The primary inspiration for this work is a recent paper of Gupta, Lee and Li [4] who found a way to "redeem" local search work for the weighted set cover for which it a priori seems doomed. A simple example which shows that local search does not work is the following. We have a set system in which the ground set has n elements. We have one set of weight 1 which covers all ground set elements, and we have n sets of weight $\epsilon \ll 1/n$ each of which covers a distinct element of the ground set. In this case, the optimal solution consists of the sets with weight ϵ . However, if we start with the set of weight 1 as our initial feasible solution, we cannot reduce the objective function (the sum of the weights of sets in the solution) by making a *small* change. The only way to reduce the

objective function is to remove the set of weight 1 from the solution and add all n sets of weight ϵ to the solution. Depending on how small ϵ is, the locally optimal solution consisting of the set of weight 1 can be arbitrarily bad compared to the optimal solution. This apparent obstacle is bypassed in [4] via *non-oblivious local search* (NOLS) (introduced in [19], see also [20], [21]). The idea is to minimize a *potential function* different from the objective function. Let X denote the ground set and \mathcal{S} denote the set of subsets of X in the given instance. At any point in time, the algorithm of [4] maintains a feasible solution $\mathcal{F} \subseteq \mathcal{S}$ along with a mapping $c : X \mapsto \mathcal{F}$ that maps each element $x \in X$ to a set $F \in \mathcal{F}$ that covers x . The potential function the algorithm seeks to minimize is the Rosenthal potential defined as:

$$\phi(\mathcal{F}, c) = \sum_{F \in \mathcal{F}} w(F) \cdot H(|c^{-1}(F)|),$$

where $w(F)$ is the weight associated with the set F , $H(m)$ is the m^{th} Harmonic number $1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{m}$ and $c^{-1}(F) = \{x \in X : c(x) = F\}$. The algorithm starts with any feasible solution \mathcal{F} and any valid mapping c . In any local move, it adds a set $S \in \mathcal{S}$ that is not in the optimal solution and modifies the mapping c as follows: for every $x \in S$, we change $c(x)$ to S . Any set F in the current solution for which $c^{-1}(F)$ becomes empty as result of this, is dropped from the current solution. A relatively simple analysis then shows that the solution obtained by repeatedly applying local moves that reduce the potential until no such move is possible, has weight at most H_k times the weight of the optimal solution. Here, k is the size of the largest set in \mathcal{S} . Note that the *size* of the local moves in this algorithm is just 1 - we only add *one* set to the current solution at a time and remove the sets made redundant by it. Local moves of larger size are also considered in the paper for lower order improvements in the approximation factor. From our point of view, the key takeaway is that local search can be "redeemed" by the use of a potential function. Since local search is a powerful heuristic frequently used in practical applications, we believe that incorporating this idea into those applications will lead to practical gains. We also believe that instead of fixing a particular potential function, learning it from data could be more effective. However, we don't pursue that avenue in this paper.

IV. OUR GENETIC ALGORITHM

Our algorithm follows the standard framework of genetic algorithms. The key components that are distinct from standard genetic algorithms are the following.

- *Fitness function.* We use the following slightly modified form of the Rosenthal potential as the fitness function:

$$\Phi(\mathcal{F}, c) = \sum_{F \in \mathcal{F}} w(F) \cdot H^\alpha(|c^{-1}(F)|),$$

where $\alpha \in [0, 1]$ is a parameter.

The parameter α used in the modified Rosenthal potential function is initially non-zero and is slowly decreased to 0. Note that when $\alpha = 1$, the objective function is the Rosenthal potential and when $\alpha = 0$, it is simply the sum of the weights of the sets in \mathcal{F} , which is the original objective of the weighted set cover problem. In this way, our algorithm initially avoids getting stuck in a local minima, but later focuses only on moves that improve upon the original objective. We note that this is somewhat similar to the idea of simulated annealing. For some instances, we even start the algorithm with an $\alpha > 1$.

One difficulty with computing Φ is that in addition to a feasible solution \mathcal{F} , it requires the map c . One way to avoid this is to define c implicitly as the mapping that minimizes the potential. However, finding such a mapping is a non-trivial optimization problem. Instead, we choose a mapping c greedily as follows. We first map each ground set element to one of the sets covering it uniformly at random. We then do one or more rounds of the following: we go over the elements one by one and find the optimal set it should be mapped to while considering the mapping of all other ground set elements fixed. Our experiments show that typically just one round suffices to obtain a good mapping. In future references to the potential function, we will not explicitly define c and assume that it is chosen by the greedy algorithm.

- *Crossover operator.* We combine two feasible solutions as follows. We start with the union U of two solutions which itself is a feasible solution and apply the following *minimalization* operation. This operation considers the sets in a feasible solution in some order and removes the set currently being considered from the solution iff the remaining sets still form a feasible solution. For the crossover operation, the order in which the sets in U are considered is determined by the contributions of the sets to the Rosenthal potential (with c computed greedily as described earlier).
- *Initial population.* We considered two approaches for creating an initial population of feasible solutions.

Random Minimalization: This approach starts with feasible solution consisting of all sets and then applies the minimalization operation while processing the sets in a random order.

Probabilistic Greedy: Recall that the greedy algorithm for set cover starts with an empty (infeasible)

solution and add sets one by one until a feasible solution is obtained. The next set S added is the one that is the most cost effective i.e., it minimizes $w(S)/\nu(S)$ where $w(S)$ is the weight of the set S and $\nu(S)$ is the number of elements that are covered by S but not by previously added sets. We make one slight modification to this in order to obtain a large number of solutions. Instead of always adding the most cost effective set, we add any set S with probability proportional to its cost effectiveness $w(S)/\nu(S)$.

The remaining details of our genetic algorithm are as follows:

- *Selection operator.* We use a 2-way tournament selection operator. For this, we choose two random pairs from the current population perform a crossover on the pair obtained by choosing the fitter individual from each random pair. This process is repeated until the required number of individuals are created for the next generation. We refer the readers to an extensive survey by Goldberg and Deb [22] for a comparative analysis of different selection operators. In addition, a fraction of the top (elite) individuals are passed from one generation to the next without any crossover operation.
- *Mutation operator.* In the mutation operator, we simply add to the current solution a few randomly selected sets that are not already part of it.

V. EXPERIMENTAL ANALYSIS

A. Dataset description

- We use the OR-library weighted SCP instances [6], [23], [17]. This dataset is a collection of test data sets for a variety of Operations Research (OR) problems and is divided into two sets OR-small and OR-large according to their size. We consider the OR-large instances in this paper.
- The second benchmark (Rail)¹ contains real-world weighted SCP instances that arise from an application in Italian railways.
- The third benchmark (STS) [24] is obtained from Steiner triple systems and consists of unweighted (or equivalently, uniformly weighted) SCP instances.

B. Experimental Set up

Next, we describe our experimental set up. Our code is implemented in C++11 and compiled using g++ with optimization flags -O3. All running times are measured on a server with an AMD EPYC 7281 16-Core Processor with 32 threads. Each core has a boost speed of 2.7 GHz.

¹<http://people.brunel.ac.uk/~mastijb/jeb/orlib/files/>

The server has a total shared memory of 96 GB and a total L3 cache of 32 MB.

For the baseline NuSC, we used the publicly available implementation². For the genetic algorithm for set cover from Beasley and Chu [14], we couldn't find any public implementation, so we implemented it ourselves based on the description in their paper.

In our genetic algorithm, the population size is set to 100 for all instances, and the number of genetic iterations is 500. The initial value of α is 1 and in each iteration, the α decays by a multiplicative factor of 0.99. We transfer 20% of elite individuals directly to the next generation without crossover operation. The mutation probability is set to 0.2 and the crossover probability is 0.8. For NuSC, we took the default values from their publicly available implementation that were tuned for the benchmark instances.

All algorithms have been run for 10 independent runs and the best and average objective function is recorded in the tables.

C. Performance on Benchmark Instances

Tables I, II and III show the comparison between the different algorithms on the weighted, unweighted and the Rail instances, respectively. Here, ILP refers to solving the integer linear programming formulation of weighted set cover using the Gurobi solver. The "obj" column of the ILP gives the optimal objective function value for the instance. As expected, the ILP takes a very long time on larger and more complex instances. We refer to our genetic algorithm with Rosenthal potential function (with α decaying from 1 by a factor of 0.99 in each iteration) and probabilistic greedy initial population as "GA (with pot. fn.)", our genetic algorithm without the Rosenthal potential function (i.e., $\alpha = 0$) as "GA (w/o pot. fn.)" the genetic algorithm from Beasley and Chu [14] as "GA-BC" and the NuSC heuristic [16] as "NuSC".

From Table I, we observe that our genetic algorithm gives near-optimal solutions on weighted instances. In fact, we get optimal solutions on these instances except for instances scpnrg1 where we get 177 instead of 176 and scpnrg3 where we get 168 instead of 166. Furthermore, we find that there is very little variance in the objective function and the time between the ten independent runs of our genetic algorithm, as indicated by the best and the average set cover weight obtained. We note that GA-BC also obtains near optimal solutions. However, this comes at the cost of a very high running time. This is because GA-BC requires a large number of iterations (generations) to obtain these solutions. The heuristic NuSC is able to achieve near-optimal solutions with significantly less running time. The results on

unweighted instances and Rail instances (Tables II and III) are similar, though on these larger instances, GA-BC can't get anywhere close to optimal in reasonable running time (less than 5 minutes).

D. Comparison with Genetic Algorithm without Potential Function

We observe that the high quality of the solutions from our genetic algorithm is due in large parts to the usage of Rosenthal potential function. Tables I, II and III show that if we were to optimize the weighted set cover objective directly, the objective function value of the resultant solutions would have been significantly worse. For instance, on scpnrh1 instance, our genetic algorithm with Rosenthal potential function yields an optimal solution with the objective of 63, while the genetic algorithm without it returns a solution with objective of 80 (27% away from the optimal solution).

E. Comparison with NuSC

We compare our genetic algorithm with the state of the art local search heuristic NuSC. As described in Section II, NuSC is a very recent approach (published in 2024) and is carefully engineered for weighted set cover problem benchmark instances. It has been shown to outperform a large number of existing heuristics on these benchmark instances [16].

Since our main motivation for using the Rosenthal potential was to escape local optima, we created simple synthetic instances which would require moves of large size with the standard local search as follows. We consider two collections of sets called A and B , which contain n and m sets respectively where $m \approx n/2$. There are $n+m$ sets in total. For each triple of sets, consisting of two sets from A and one set from B , we create a ground set element present in exactly those three sets. Thus, the number of ground set elements is $m \cdot \binom{n}{2}$. It can be checked both A and B are feasible solutions and B is an optimal solution to the unweighted set cover problem. Despite the large gap in sizes, if we start from A , there is no local move that improves the solution other than the move that changes A to B directly - involving a large local move (of size proportional to n). We also extend these instances by replacing each set in A by k copies of the same set for some $k > 1$. We now have $kn+m$ sets and the number of ground set elements remains $m \cdot \binom{n}{2}$.

Note that on these instances, the linear programming (LP) relaxation of the set cover integer linear program will result in all sets from the collection A getting the value $1/2$ while all sets from the collection B will get the value 0. Thus, heuristics based on LP rounding techniques or those that just return sets with non-zero LP values will fail on these instances.

²<https://github.com/chuanluocs/NuSC-Algorithm>

Table I: Comparison of different algorithms on weighted SCP instances from the OR Library

Instance	ILP		Our GA (with pot. fn.)		GA (w/o pot. fn.)		GA-BC		NuSc	
	obj	time	min-obj (avg)	time	min-obj (avg)	time	min-obj (avg)	time	min-obj (avg)	time
scpnrg1	176	791.9	177(177.6)	418.54	195(195.7)	387.5	178(180.1)	2124.2	176(176)	0.08
scpnrg2	154	318.8	154(155.3)	417.67	165(167.5)	402.4	158(159)	2613.4	154(154)	0.14
scpnrg3	166	3510.5	168(168.8)	419.99	183(184.5)	378.8	168(169.2)	2921.7	166(166)	2.16
scpnrg4	168	2339.2	168(169.8)	420.10	189(192.1)	398.2	170(171.6)	2431.3	168(168)	99.28
scpnrg5	168	8099.3	168(169)	420.10	188(188.8)	365.7	170(171)	2812.2	168(168)	2.55
scpnrh1	63	300259.5	63(63.7)	498.95	80(80.3)	335.9	64(64.8)	3701.3	63(63)	3.49
scpnrh2	63	121093.7	63(63.9)	499.42	78(78.9)	387.3	64(64.8)	3503.3	63(63)	0.44
scpnrh3	59	12897.6	59(60)	494.22	65(67)	412.4	60(61)	3723.9	59(59)	1.43
scpnrh4	58	26613.8	58(58.8)	491.69	65(66.5)	387.2	59(59.9)	3402.3	58(58)	0.91
scpnrh5	55	33843.5	55(55.8)	477.54	62(63)	404.2	55(58.2)	3801.1	55(55)	0.5

Table II: Comparison of different algorithms on unweighted SCP instances from the OR library and the Steiner triple system (sts) instances

Instance	Size	ILP		Our GA (with pot. fn.)		GA (w/o pot. fn.)		GA-BC		NuSc	
	row×column	min-obj	time	min-obj (avg)	time	bj	time	min-obj (avg)	time	min-obj (avg)	time
scpcrl0	511×210	25	3.2	25(26.7)	15.2	47(48.3)	20.9	25(26.7)	28.9	25(25)	0.0
scpcrl1	1023×330	23	88.6	25(25.2)	230.8	54(57)	219.1	25(26.2)	35.6	23(23)	0.04
scpcrl2	2047×495	23	1304.0	26(26.2)	610.3	44(45.6)	596.8	26(26.5)	39.3	23(23)	0.35
scpcrl3	4095×715	23	19772.7	25(25.2)	1211.3	43(44.5)	1102.9	25(26.5)	46.7	23(23)	0.61
scpcyc06	240×192	60	1002.4	62(62.1)	21.0	70(74.2)	21.1	66(68.0)	61.4	60(60)	0.0
scpcyc07	672×448	144	1002.1	148(148.8)	87.1	161(165.5)	83.2	160(160.8)	66.3	144	0.03
scpcyc08	1792×1024	342	1003.3	360(362.1)	283.4	392(398.1)	291.3	402(408.8)	79.5	344(344)	30.52
scpcyc09	4068×2304	772	1001.5	850(885)	1413.1	892(900.8)	1398.3	911(917.3)	115.7	780(780)	796.36
scpcyc10	11520×5120	1798	1000.9	1992(1996.8)	3098.3	2087(2092.6)	2871.3	2108	217.6	1794	340.15
scpcyc11	28160×11264	3968	1004.3	4104(4108.4)	4873.4	4687(4693.3)	4242.4	4398(4404.5)	783.4	3968(3968)	288.08
sts135	3015×135	103	–	106(106.5)	117.1	118(119.6)	129.6	106(106.5)	126.8	103(103)	158.67
sts243	9801×243	198	–	204(204.5)	508.1	227(227.8)	563.1	206(206.5)	618.3	198(198)	0.0
sts405	27270×405	335	–	344(345.7)	832.3	379(282.4)	862.4	350(250.5)	922.4	336(336.5)	10.33
sts729	88452×729	617	9142.1	628(630.1)	4974.9	687(690.3)	4871.3	652(653.1)	2118.2	617(617)	52.19

Table III: Comparison of different instances on Rail instances

Instance	Size	ILP		Our GA (with pot. fn.)		GA (w/o pot. fn.)		NuSc	
	row×column	min-obj	time	min-obj (avg)	time	min-obj (avg)	time	min-obj (avg)	time
rail507	507×63009	174	104.8	182(182.5)	873.1	297(299.4)	814.6	174(174)	528.79
rail516	516×47311	182	56.7	189(189.3)	1841.4	286(288.2)	1791.5	182(182)	2.95
rail582	582×55515	211	89.2	216(216.8)	2487.3	298(299.1)	2413.3	211(211)	50.46
rail2536	2536×1081841	689	6277.6	726(727)	4481.5	954(958.5)	4173.9	699(699.4)	185.51
rail2586	2586×920683	951	11092.8	1108(1109.2)	6723.2	1985(1985.6)	6034.3	960(961.5)	866.2

Table IV: Comparison of our Genetic algorithm and NuSc on synthetic instances. Ins(m,n,k) is an instance with m sets in collections B , n sets in collection A and k copies of each set in A .

Instance	Size	ILP		Our GA (with pot. fn.)		GA (w/o pot. fn.)		NuSc	
	row×column	obj	time	min-obj (avg)	time	min-obj (avg)	time	min-obj (avg)	time
Ins(7,13,1)	546×20	7	0.03	7(7)	0.2	7(7)	0.2	7(7)	205.9
Ins(11,20,1)	2090×31	11	0.21	11(11)	1.3	19(19)	45.5	11(11)	403.5
Ins(16,30,1)	6960×46	16	1.06	16(16)	2.2	29(29)	169.9	16(16)	2540.3
Ins(51,100,1)	252450×151	51	789.9	51(51)	180.6	99(99)	3000.0	99(99)	3000
Ins(7,13,2)	546×33	7	0.05	7(7)	1.2	12(12)	13.0	7(7)	981.0
Ins(11,20,2)	2090×51	11	0.34	11(11)	3.1	19(19)	49.2	19(19)	3000
Ins(16,30,2)	6960×76	16	1.60	16(16)	17.8	29(29)	187.2	29(29)	3000
Ins(51,100,2)	252450×251	51	615.3	51(51)	214.2	99(99)	3000.0	99(99)	3000

As seen in Table IV, our genetic algorithm (with $\alpha = 4$ and decay factor of 0.98) always obtains the optimal solution on these instances. In contrast, NuSC is often quite far from the optimal solution even after the time-out of 3000 seconds. This is particularly true for instance that are large or instances that have multiple copies of A . For instance, on Ins(51,100,1), NuSC obtains a solution of 99 while the optimal solution has the objective value of 51. For these instances, NuSC is reporting the sets from the collection A in the returned solution while the actual optimal solution consists of sets from the collection B . The results with 4, 6 and 8 copies of sets in A are similar to that of 2 copies.

Again, we note that without the Rosenthal potential function, our genetic algorithm would have returned solutions similar to the NuSC heuristic or even worse. Thus, we conclude that the main reason why our genetic algorithm is able to do local moves of larger size on these instances is the usage of Rosenthal potential function.

VI. CONCLUSION

Our experiments indicate that incorporating the Rosenthal potential has a significant impact on the quality of the solution obtained. While our genetic algorithm is able to match the quality of the solution obtained for benchmark instances, NuSC is faster on those instances. The difference in running times stems primarily from the large populations sizes maintained in a genetic algorithm. On the other hand, the advantage of the our genetic algorithm with modified Rosenthal potential is clear when we have instances requiring large local moves. For instance, on synthetic instances requiring large local moves, our algorithm gets optimal solutions in time that is often two orders of magnitude less than the state-of-the-art heuristic NuSC (which typically fails to obtain the optimal solution even after a long time). It is an interesting challenge to combine the advantages of both the algorithms. Another interesting direction is to find other applications where modifying the objective function improves the practical performance of local search.

REFERENCES

- [1] V. Chvatal, "A greedy heuristic for the set-covering problem." *Math. Oper. Res.*, vol. 4, no. 3, pp. 233–235, 1979. doi: <https://doi.org/10.1287/MOOR.4.3.233>
- [2] I. Dinur and D. Steurer, "Analytical approach to parallel repetition," in *STOC*. ACM, 2014. doi: <https://doi.org/10.1145/2591796.2591884> p. 624–633.
- [3] K. Clarkson and K. Varadarajan, "Improved approximation algorithms for geometric set cover," *Discrete Computational Geometry*, vol. 37, pp. 43–58, 2007. doi: <https://doi.org/10.1007/S00454-006-1273-8>
- [4] A. Gupta, E. Lee, and J. Li, "A local search-based approach for set covering," in *SOSA*. SIAM, 2023. doi: <https://doi.org/10.1137/1.9781611977585.CH1> pp. 1–11.
- [5] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *Int. Jour. of Game Theory*, vol. 2, pp. 65–67, 1973.
- [6] J. E. Beasley, "An algorithm for set covering problem," *European Journal of Operational Research*, vol. 31, no. 1, pp. 85–93, 1987. doi: [https://doi.org/10.1016/0377-2217\(87\)90141-X](https://doi.org/10.1016/0377-2217(87)90141-X)
- [7] N. Bansal, A. Caprara, and M. Sviridenko, "A new approximation method for set covering problems, with applications to multidimensional bin packing," *SIAM J. Comput.*, vol. 39, pp. 1256–1278, 2009. doi: <https://doi.org/10.1137/080736831>
- [8] J. Bautista and J. Pereira, "A grasp algorithm to solve the unicast set covering problem," *Computers & Operations Research*, vol. 34, no. 10, pp. 3162–3173, 2007. doi: <https://doi.org/10.1016/j.cor.2005.11.026>
- [9] Y. Wang, S. Pan, S. Al-Shihabi, J. Zhou, N. Yang, and M. Yin, "An improved configuration checking-based algorithm for the unicast set covering problem," *EJOR*, vol. 294, no. 2, pp. 476–491, 2021. doi: <https://doi.org/10.1016/j.ejor.2021.02.015>
- [10] Z. Naji-Azimi, P. Toth, and L. Galli, "An electromagnetism metaheuristic for the unicast set covering problem," *EJOR*, vol. 205, no. 2, pp. 290–300, 2010. doi: <https://doi.org/10.1016/j.ejor.2010.01.035>
- [11] C. Gao, X. Yao, T. Weise, and J. Li, "An efficient local search heuristic with row weighting for the unicast set covering problem," *EJOR*, vol. 246, no. 3, pp. 750–761, 2015. doi: <https://doi.org/10.1016/j.ejor.2015.05.038>
- [12] Z. Lei and S. Cai, "Solving set cover and dominating set via maximum satisfiability," in *EAAI*. AAAI Press, 2020. doi: <https://doi.org/10.1609/AAAI.V34I02.5517> pp. 1569–1576.
- [13] M. J. Brusco, L. W. Jacobs, and G. M. Thompson, "A morphing procedure to supplement a simulated annealing heuristic for cost- and coverage-correlated set-covering problems," *Ann. Oper. Res.*, vol. 86, pp. 611–627, 1999. doi: <https://doi.org/10.1023/A%3A1018900128545>
- [14] J. Beasley and P. Chu, "A genetic algorithm for the set covering problem," *EJOR*, vol. 94, no. 2, pp. 392–404, 1996. doi: [https://doi.org/10.1016/0377-2217\(95\)00159-X](https://doi.org/10.1016/0377-2217(95)00159-X)
- [15] B. Crawford, R. Soto, R. Cuesta, and F. Paredes, "Application of the artificial bee colony algorithm for solving the set covering problem," *Scientific World Journal*, 2014. doi: <https://doi.org/10.1155/2014/189164>
- [16] C. Luo, W. Xing, S. Cai, and C. Hu, "Nusc: An effective local search algorithm for solving the set covering problem," *IEEE Trans. Cybern.*, vol. 54, no. 3, pp. 1403–1416, 2024. doi: <https://doi.org/10.1109/TCYB.2022.3199147>
- [17] T. Grossman and A. Wool, "Computational experience with approximation algorithms for the set covering problem," *EJOR*, vol. 101, no. 1, pp. 81–92, 1997. doi: [https://doi.org/10.1016/S0377-2217\(96\)00161-0](https://doi.org/10.1016/S0377-2217(96)00161-0)
- [18] S. Katoch, S. Chauhan, and V. Kumar, "A review on genetic algorithm: past, present, and future," *Multimed Tools Appl*, vol. 80, p. 8091–8126, 2021. doi: <https://doi.org/10.1007/s11042-020-10139-6>
- [19] S. Khanna, R. Motwani, M. Sudan, and U. Vazirani, "On syntactic versus computational views of approximability," *SIAM Journal on Computing*, vol. 28, no. 1, pp. 164–191, 1998.
- [20] Y. Filmus and J. Ward, "Monotone submodular maximization over a matroid via non-oblivious local search," *SIAM Journal on Computing*, vol. 43, no. 2, pp. 514–542, 2014.
- [21] V. Cohen-Addad, A. Gupta, L. Hu, H. Oh, and D. Saulpic, "An improved local search algorithm for k-median," in *SODA*. SIAM, 2022. doi: [10.1137/1.9781611977073.65](https://doi.org/10.1137/1.9781611977073.65) pp. 1556–1612.
- [22] D. E. Goldberg and K. Deb, "A comparative analysis of selection schemes used in genetic algorithms," in *First Workshop on Foundations of Genetic Algorithms*, vol. 1. Elsevier, 1991. doi: <https://doi.org/10.1016/B978-0-08-050684-5.50008-2> pp. 69–93.
- [23] J. E. Beasley, "A lagrangian heuristic for set-covering problems," *Naval Research Logistics*, vol. 37, pp. 151–164, 1990. doi: <https://doi.org/10.1002/1520-6750>
- [24] D. Fulkerson, G. L. Nemhauser, and L. Trotter, "Two computationally difficult set covering problems that arise in computing the 1-width of incidence matrices of steiner triple systems," in *Approaches to integer programming*, 1974. doi: <https://doi.org/10.1007/BFb0120689> pp. 72–81.

Efficiency and Reliability of Avalanche Consensus Protocol in Vehicular Communication Networks

Saeed Ullah

Quaid I Azam University
Islamabad Pakistan

Email: saeed.ullah@ele.qau.edu.pk

Zaib Ullah

Università Telematica Giustino Fortunato
Benevento Italy

Email: z.ullah@unifortunato.eu

Abdullah Waqas

National University of Technology
Islamabad Pakistan

Email: abdullah@nutech.edu.pk

Abstract—In vehicular communication networks, centralized systems face significant security challenges, including privacy preservation, secure authentication, threats from compromised authorities, latency, and throughput. We propose a blockchain-based system that decentralizes control, enhances throughput, and optimizes latency. By leveraging the Avalanche consensus protocol, our solution assures efficient, secure, and robust communication within vehicular networks, mitigating risks associated with centralized control. Our proposed system achieves a substantial throughput, with the Practical Byzantine Fault Tolerance (PBFT) protocol registering 12.8 transactions per second (TPS), and the Avalanche protocol demonstrates an impressive 1007 TPS for 100 validators. Regarding the delay, PBFT experiences 6.61 seconds, whereas Avalanche protocol achieves a remarkably low delay of just one millisecond, both with 100 validators. These findings highlight the superiority of our proposed system in terms of low latency, and enhanced transaction throughput, essential for future vehicular communication systems.

I. INTRODUCTION

IN RECENT years, blockchain technology has emerged as a transformative solution for enhancing security and decentralization across various domains. Originally conceived for cryptocurrencies, blockchain's immutable and distributed ledger system offers robust security features, making it an attractive option for applications requiring secure and transparent data handling. The effective use of blockchain technology plays a crucial role in building smart cities [1]. Additionally, blockchain technology is vital in monitoring systems that assist elderly and impaired patients in adhering to their medication regimens at home. By leveraging blockchain, these systems ensure that all patient-related activities are securely logged and verified, providing a reliable and secure framework for patient care management [2].

In the realm of vehicular communication, where issues such as privacy preservation [3], secure authentication, and protection against insider threats are paramount, blockchain presents a promising approach to address these challenges. By decentralizing control and eliminating single points of failure, blockchain can significantly enhance the resilience and trustworthiness of vehicular communication networks.

Numerous researchers have explored the application of blockchain in vehicular communication to mitigate security

vulnerabilities. For instance, studies have proposed the use of blockchain for secure vehicular identity management and authentication, aiming to prevent unauthorized access and ensure that only legitimate vehicles can communicate within the network. Additionally, blockchain-based solutions have been developed to safeguard data integrity and privacy, ensuring that sensitive information shared between vehicles remains confidential and tamper-proof. Despite these advancements, the centralized nature of some proposed solutions still poses risks, as centralized authorities can be single points of failure or targets for attacks [4].

Moreover, the challenge of achieving high transactions per second (TPS) in vehicular communication networks has been a significant hurdle. Traditional blockchain systems, such as Bitcoin and Ethereum, are often criticized for their limited scalability and low TPS, which are insufficient for the high-frequency data exchanges required in vehicular networks. Various consensus mechanisms, including Proof of Work (PoW) and Proof of Stake (PoS), have been explored to improve scalability, but they often fall short of meeting the stringent performance requirements of vehicular communication. To address these issues, we introduce a novel consensus algorithm based on the Avalanche protocol, an innovative and scalable method designed to enhance TPS while maintaining high security standards. Our approach demonstrates the potential to achieve 1007 TPS with 100 validators, providing a robust and efficient framework for secure and high-performance vehicular communication networks.

The rest of the paper is organized as follows: **Section 2** dives into the technical foundation, covering blockchain's secure ledgers and consensus mechanisms. **Section 3** introduces the core functionalities of blockchain technology and provides the groundwork for understanding our proposed solution. **Section 4** defines the system model for secure vehicular communication. **Section 5** provides analyses of the attained results. **Section 6** explores the challenges and promising future research directions. Finally, **Section 7** concludes the article by summarizing the key findings.

II. TECHNICAL BACKGROUND

Blockchain technology is a distributed ledger system that securely, transparently, and immutably records transactions. Each block in the chain contains a list of transactions, a timestamp, and a cryptographic hash of the preceding block, creating an alteration-resistant record. This decentralized architecture eliminates single points of control, fostering security and trust among participants [5].

In a blockchain network, participants initiate and broadcast transactions across the network. Validators, also known as miners in some systems, collect these transactions into blocks. Before adding a block to the blockchain, a consensus mechanism ensures its validity by achieving network-wide agreement on the included transactions. This process prevents double-spending and safeguards against fraudulent activities [5].

The consensus mechanism is a fundamental component of blockchain technology, as it guarantees the integrity and consistency of the distributed ledger. A variety of consensus mechanisms exist, each with its strengths and weaknesses. Proof of Work (PoW) [6], used by Bitcoin blockchain, requires validators to solve complex cryptographic puzzles, which ensures security but is energy-intensive and slow [5], [7]. Proof of Stake (PoS) [6], on the other hand, selects validators based on their stake in the network, which is more energy-efficient but can lead to centralization if a few participants hold large stakes [8]. The consensus mechanism validates transactions determines how the blockchain grows and ensures that all copies of the distributed ledger are synchronized across the network [9], [10], [11].

In vehicular communication, where rapid data exchange and high security are paramount, traditional consensus mechanisms often prove inadequate. The high latency and low throughput associated with PoW and PoS make them unsuitable for the real-time demands of vehicular networks. These limitations have prompted researchers to explore alternative consensus algorithms that can provide higher TPS while maintaining robust security standards. One such innovative approach is the Avalanche consensus protocol, which uses repeated sampling and confidence levels to achieve rapid consensus without the need for extensive computational resources. By leveraging the Avalanche protocol, it is possible to significantly improve the TPS in vehicular communication networks, addressing one of the major bottlenecks in existing blockchain solutions.

For instance, studies have proposed blockchain-based identity management systems for vehicles. These systems leverage the blockchain's immutable ledger to store and verify vehicle identities, enhancing trust and preventing unauthorized access within the network [12], [13], [14], [15], [16], [17]. Additionally, blockchain's ability to provide a tamper-proof record of transactions makes it well-suited for preserving the privacy and integrity of data exchanged between vehicles [18], [19], [20]. However, these solutions still face challenges related to scalability and performance, highlighting the need for more efficient consensus mechanisms like the Avalanche protocol discussed earlier.

Integrating blockchain with machine learning in vehicular communication networks can significantly enhance security, efficiency, and data analysis capabilities. Blockchain can provide a secure and immutable data storage solution, ensuring the integrity and privacy of the data collected from vehicles. Machine learning algorithms can then analyze this data to predict traffic patterns, optimize routing, and enhance autonomous driving capabilities. For instance, secure data sharing through blockchain can support federated learning models, where decentralized data can be collaboratively processed without compromising privacy [21], [22]. This integration creates a robust framework for real-time decision-making and continuous improvement of vehicular systems.

The Avalanche consensus protocol [23] represents a promising advancement in this field. By enabling high TPS and reducing latency, it addresses the critical performance issues that have hindered the adoption of blockchain in vehicular communication. The proposed system, which achieves 1007 TPS with 100 validators, demonstrates the potential of Avalanche to meet the strict requirements of vehicular networks, providing a secure, decentralized, and high-performance solution for future applications.

III. INTRODUCTION TO BLOCKCHAIN AND CONSENSUS MECHANISMS

In the following, we briefly explore the fundamentals of blockchain technology and consensus mechanisms.

A. Blockchain

Blockchain technology, originally conceived as the foundation for Bitcoin cryptocurrency, has rapidly transformed into a versatile solution across diverse industries, including vehicular communication networks. At its core, it functions like a public record book, not controlled by a single entity but maintained by a distributed network of computers. This distributed ledger, called a blockchain, consists of interconnected blocks as shown in Fig. 1 [24], each containing a list of transactions. These transactions can represent anything of value, such as financial transfers or ownership records.

The core strength of blockchain lies in its decentralization, eliminating the need for a central authority and making it resistant to manipulation. Additionally, the immutability of transactions, meaning they cannot be altered once recorded, ensures data integrity and transparency. Cryptographic techniques further secure the blockchain, offering a robust shield against tampering and fraud.

In the context of vehicular communication, blockchain offers significant advantages. It can facilitate secure and reliable data exchange among vehicles and infrastructure, addressing critical concerns like unauthorized access and data manipulation. The tamper-proof nature of blockchain also helps preserve the privacy of data exchanged between vehicles. Furthermore, the decentralized architecture enhances the resilience of vehicular communication networks by eliminating single points of failure. However, existing solutions leveraging blockchain in this domain still face challenges related to



Fig. 1. Blockchain Technology

scalability and performance. This is where innovative consensus mechanisms like the Avalanche protocol come into play, aiming to address these limitations by achieving high transaction throughput and low latency, paving the way for secure and efficient vehicular communication networks of the future.

B. Consensus Algorithm

The consensus algorithm is a critical component of blockchain technology, ensuring that all distributed nodes agree on the state of the ledger. The primary goal is to enable each node to verify block generation in a decentralized manner, typically by selecting one validator or miner per round to add a new block. Three common consensus algorithms are Proof of Work (PoW), Proof of Stake (PoS), and PBFT. In PoW, validators compete by solving complex cryptographic puzzles, while PoS relies on validators who have staked cryptocurrency to participate in block creation.

1) *Practical Byzantine Fault Tolerance*: The PBFT consensus algorithm employs a two-thirds majority voting method, making it suitable for private blockchains where trusted participants are known. It's commonly used in systems like IBM Hyperledger Fabric. As illustrated in Fig 2, the PBFT process involves distinct stages: REQUEST, PRE-PREPARE, PREPARE, COMMIT, and REPLY [25]. To understand how Byzantine faults are tolerated, here's a breakdown of the PBFT process:

1) Request:

- The client initiates the process by broadcasting a transaction to all nodes in the network.

2) Pre-Prepare:

- The primary node, selected through a pre-defined process or upon receiving the first client message, assembles a block containing the received transaction(s).
- This block is then broadcast to all replica nodes in the network.

3) Prepare:

- Upon receiving the pre-prepared block, each replica node verifies two things:
 - a) Whether it has received the same pre-prepared block from the primary node.
 - b) Whether the transactions and values within the block are valid according to the system's rules.
- If both conditions are met, the replica node broadcasts a "prepare" message for the block.

4) Commit:

- Once a replica node receives "prepare" messages from more than two-thirds of the other nodes (including the primary node), it can be confident that the block is legitimate.
- This threshold ensures resistance to Byzantine faults, where nodes might malfunction or provide malicious information.
- At this stage, the replica node broadcasts a "commit" message for the block.

5) Reply:

- Finally, all nodes, including the primary and replica nodes, send a reply message to the client, indicating successful transaction processing or any potential errors encountered.




Fig. 2. PBFT Consensus Protocol

2) *Avalanche Consensus Protocol*: The Avalanche consensus protocol is a unique consensus mechanism designed to achieve high throughput, low latency, and robust security. Here is an overview of the Avalanche consensus algorithm:

1) Initialization:

- Nodes in the network are initialized and are aware of each other. Each node starts with an initial set of transactions it considers valid.

2) Transaction Propagation:

- A node receives a transaction and propagates it to a small, randomly selected subset of nodes (neighbors).

3) Voting Process:

- Each node in the subset checks the validity of the transaction and votes to accept or reject it based on its local view.
- The voting process involves repeated rounds where nodes query their peers to reach a consensus.

4) Sampling and Subsampling:

- In each round, nodes sample a random subset of other nodes (k nodes) and adopt the decision of the majority within this subset.
- This process continues until a majority consensus is reached consistently over several rounds.

5) Confidence Levels:

- Each node maintains a confidence level for each transaction based on the number of times the transaction is voted as valid in consecutive rounds.
- A transaction is considered finalized when the confidence level surpasses a predefined threshold.

6) Finalization:

- Once the confidence threshold is reached, the transaction is finalized and accepted by the network.

IV. SYSTEM MODEL

The proposed system model for secure and efficient vehicular communication leverages the Avalanche consensus protocol to address the inherent challenges of achieving high throughput and low latency in vehicular networks. This system model includes several key components and mechanisms designed to ensure robust security, privacy, and performance.

The vehicular network, depicted in Figure 3 as a dynamic Vehicular Ad Hoc Network (VANET) [26], comprises a large number of vehicles acting as nodes in a decentralized blockchain network. Each vehicle is equipped with a communication module capable of interacting with other vehicles and the blockchain network. These vehicles perform dual roles: they act as both users initiating transactions and validators participating in the consensus process.

To initiate a transaction, a vehicle broadcasts its message to the network. This message may include data such as vehicle identity, location, or other relevant information necessary for the intended application (e.g., traffic management, collision avoidance). The transaction is then received by neighboring vehicles, which subsequently propagate the transaction across the network.

Each transaction is grouped into a block by a validator. Unlike traditional PoW and PoS systems, the proposed model employs the Avalanche consensus protocol, which is well-suited for environments requiring rapid consensus and high transaction throughput. Validators in the Avalanche protocol operate by continuously sampling the network to determine the confidence level of each transaction. This process involves selecting a small, random subset of validators and querying their opinions on the validity of a transaction. Validators update their own state based on the responses received, and this process is repeated multiple times until a predefined confidence threshold is met.

The confidence level is a critical parameter in this model. It determines the number of repeated sampling rounds necessary to reach consensus. This parameter is set to ensure a balance between security and performance, allowing the system to maintain high throughput while providing strong guarantees against attacks and inconsistencies.

To enhance security and privacy, the system model includes mechanisms for secure authentication and identity management. Each vehicle is issued a unique cryptographic key pair, which it uses to sign its transactions. This ensures that only authorized vehicles can participate in the network and provides a means for verifying the integrity and authenticity of each transaction. Additionally, the blockchain's immutable ledger




Fig. 3. Vehicular Ad-hoc Network

serves as a tamper-proof record of all transactions, further enhancing security by making it difficult for malicious actors to alter or forge data.

Privacy preservation is addressed through the use of pseudonymous identities. Vehicles do not reveal their true identities when broadcasting transactions; instead, they use temporary pseudonyms that can be periodically changed to prevent tracking and linking of activities. This approach ensures that while the network maintains a transparent and verifiable record of transactions, the privacy of individual vehicles is preserved.

The proposed system model also includes provisions for handling compromised nodes and insider threats. The Avalanche consensus protocol's reliance on repeated sampling and confidence levels makes it resilient to Sybil attacks and other forms of manipulation. Even if a subset of validators is compromised, the probability of them consistently influencing the consensus process is minimized, ensuring the integrity and reliability of the network.

V. RESULTS

In Figure 4, the comparison of TPS between PBFT and Avalanche consensus mechanisms is illustrated. PBFT achieves a TPS of 12.8 [25], while Avalanche reaches 1007 TPS. This obvious contrast highlights Avalanche's superior scalability and efficiency in processing high volumes of transactions. The increased TPS in Avalanche can be attributed to its probabilistic consensus mechanism, which reduces communication overhead, enabling faster transaction processing compared to the deterministic approach of PBFT.

Figure 5 focuses on the delay experienced in reaching consensus. For a network with 100 validators, PBFT exhibits a delay of 6.61 seconds [25], whereas Avalanche demonstrates a remarkably low delay of just 1 millisecond. This minimal delay in Avalanche is due to its innovative consensus algorithm, which leverages repeated random sampling and metastability to achieve quick and efficient consensus. In contrast, PBFT's higher delay results from its multiple rounds of communication




Fig. 4. TPS vs Number of Vlaiadors

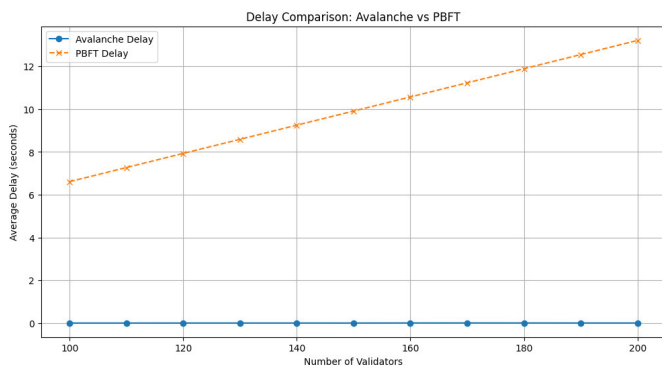


Fig. 5. Delay vs Number of Vlaiadors

and voting among validators, which are necessary to ensure fault tolerance.

The implications of these results are particularly significant for vehicle-to-vehicle (V2V) communication systems. High throughput and low latency are critical for the rapid exchange of information necessary for safety and coordination among vehicles. Avalanche's ability to handle a high number of transactions per second with minimal delay makes it a suitable candidate for such applications, ensuring that communication between vehicles is fast, reliable, and efficient. On the other hand, PBFT, with its higher delay and lower TPS, may struggle to meet the stringent requirements of V2V communication, where timely data exchange is paramount.

VI. CHALLENGES AND FUTURE RESEARCH DIRECTIONS

The Avalanche protocol offers a glimpse into a future of secure and scalable V2X communication. However, several challenges need to be addressed before widespread adoption:

- Achieving the theoretical throughput of Avalanche in real-world V2X networks with high vehicle density remains an open question.
- Ensuring user privacy through privacy-preserving consensus and data anonymization techniques is essential [27].
- Seamless integration with existing V2X infrastructure requires careful consideration to minimize disruption and ensure compatibility.

Looking towards the future, several research directions hold promise for more efficient and adaptable solutions:

- Hybrid consensus models offer potential for V2X communication but require further research for optimal integration within the V2X ecosystem [28].
- Developing lightweight blockchain clients specifically designed for resource-constrained vehicles can significantly reduce computational overhead and improve energy efficiency.
- Rigorous formal verification of the Avalanche protocol's security properties is essential for wider adoption.
- Developing standard protocols for integrating Avalanche with V2V and V2X communication systems is also essential for wider adoption.

By handling these challenges and pursuing these promising research directions, the Avalanche protocol has the potential to revolutionize V2X communication, paving the way for a secure, reliable, and efficient future for intelligent transportation systems.

VII. CONCLUSION

Blockchain technology is a cornerstone of decentralized applications, particularly in V2X communication where decentralization is paramount for security and reliability. The Avalanche protocol stands out for its high throughput and low latency, making it a strong candidate for V2X communication. This analysis revealed Avalanche's impressive performance, achieving a TPS of 1007 with a delay of just 1 millisecond for 100 validators, significantly exceeding PBFT's capabilities (TPS: 12.8, delay: 6.61 seconds). These findings highlight Avalanche's potential to effectively address the demanding requirements of modern decentralized applications compared to PBFT. However, unlocking the full potential of Avalanche in V2X networks requires further exploration. Key challenges include achieving scalability in dense real-world scenarios, designing robust incentive mechanisms for validators, and ensuring user privacy through data anonymization techniques. By addressing these challenges and pursuing these advancements, Avalanche has the potential to revolutionize V2X communication, paving the way for a secure, reliable, and efficient future for intelligent transportation systems.

ACKNOWLEDGMENT

This research has been partially funded by the Regione Campania with the "PSR Campania 2014-2022 programme, Misura 16, Tipologia d'intervento 16.1.2." Project: "EVOOLIO - L'Evoluzione dell'Olio EVO Sannita tracciato con la Blockchain.

REFERENCES

- [1] Z. Ullah, M. Naeem, A. Coronato, P. Ribino, and G. De Pietro, "Blockchain applications in sustainable smart cities," *Sustainable Cities and Society*, p. 104697, 2023.
- [2] G. Paragliola, A. Coronato, M. Naeem, and G. De Pietro, "A reinforcement learning-based approach for the risk management of e-health environments: A case study," in *2018 14th international conference on signal-image technology & internet-based systems (SITIS)*. IEEE, 2018, pp. 711-716.

- [3] A. Coronato, G. de Pietro, and G. Paragliola, "A monitoring system enhanced by means of situation-awareness for cognitive impaired people," in *BodyNets '13: Proceedings of the 8th International Conference on Body Area Networks*. Brussels, BEL: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Sep. 2013, pp. 124–127.
- [4] M. Jamal, Z. Ullah, M. Naeem, M. Abbas, and A. Coronato, "A hybrid multi-agent reinforcement learning approach for spectrum sharing in vehicular networks," *Future Internet*, vol. 16, no. 5, p. 152, 2024.
- [5] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [6] S. Yan, "Analysis on blockchain consensus mechanism based on proof of work and proof of stake," in *2022 International Conference on Data Analytics, Computing and Artificial Intelligence (ICDACAI)*, 2022, pp. 464–467.
- [7] S. S. Mahdi, Z. Ullah, G. Battineni, M. G. Babar, and U. Daoud, "The telehealth chain: a framework for secure and transparent telemedicine transactions on the blockchain," *Irish Journal of Medical Science (1971-)*, pp. 1–9, 2024.
- [8] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," *white paper*, vol. 3, no. 37, pp. 2–1, 2014.
- [9] M. Naeem, S. Bashir, M. U. Khan, and A. A. Syed, "Performance comparison of scheduling algorithms for mu-mimo systems," in *2016 13th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, 2016, pp. 601–606.
- [10] S. Islam, M. J. Islam, M. Hossain, S. Noor, K.-S. Kwak, and S. M. R. Islam, "A survey on consensus algorithms in blockchain-based applications: Architecture, taxonomy, and operational issues," *IEEE Access*, vol. 11, pp. 39 066–39 082, 2023.
- [11] M. Jamal, Z. Ullah, and M. Abbas, "Self-adapted resource allocation in v2x communication," in *Workshop Proceedings of the 19th International Conference on Intelligent Environments (IE2023)*, vol. 32. IOS Press, 2023, p. 104.
- [12] Z. Lu, Q. Wang, G. Qu, H. Zhang, and Z. Liu, "A blockchain-based privacy-preserving authentication scheme for vanets," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 12, pp. 2792–2801, 2019.
- [13] M. Naeem, G. Paragliola, A. Coronato, and G. De Pietro, "A cnn based monitoring system to minimize medication errors during treatment process at home," in *Proceedings of the 3rd International Conference on Applications of Intelligent Systems*, 2020, pp. 1–5.
- [14] C. Wang, J. Shen, J.-F. Lai, and J. Liu, "B-tsca: Blockchain assisted trustworthiness scalable computation for v2i authentication in vanets," *IEEE Transactions on Emerging Topics in Computing*, vol. 9, no. 3, pp. 1386–1396, 2021.
- [15] D. Gabay, K. Akkaya, and M. Cebe, "Privacy-preserving authentication scheme for connected electric vehicles using blockchain and zero knowledge proofs," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 6, pp. 5760–5772, 2020.
- [16] Q. Feng, D. He, S. Zeadally, and K. Liang, "Bpas: Blockchain-assisted privacy-preserving authentication system for vehicular ad hoc networks," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 6, pp. 4146–4155, 2020.
- [17] C. Lin, D. He, X. Huang, N. Kumar, and K.-K. R. Choo, "Beppa: A blockchain-based conditional privacy-preserving authentication protocol for vehicular ad hoc networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 12, pp. 7408–7420, 2021.
- [18] A. Alharthi, Q. Ni, and R. Jiang, "A privacy-preservation framework based on biometrics blockchain (bbc) to prevent attacks in vanet," *IEEE Access*, vol. 9, pp. 87 299–87 309, 2021.
- [19] G. Paragliola and M. Naeem, "Risk management for nuclear medical department using reinforcement learning algorithms," *Journal of Reliable Intelligent Environments*, vol. 5, pp. 105–113, 2019.
- [20] D. Zheng, C. Jing, R. Guo, S. Gao, and L. Wang, "A traceable blockchain-based access authentication system with privacy preservation in vanets," *IEEE Access*, vol. 7, pp. 117 716–117 726, 2019.
- [21] A. R. Javed, M. M. Hassan, F. Shahzad, W. Ahmed, S. Singh, T. Baker, and T. R. Gadekallu, "Integration of blockchain technology and federated learning in vehicular (iot) networks: A comprehensive survey," *Sensors*, vol. 22, no. 12, p. 4394, Jun 2022.
- [22] M. Dibaei, X. Zheng, Y. Xia, and X. Xu, "Investigating the prospect of leveraging blockchain and machine learning to secure vehicular networks: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. PP, no. 99, pp. 1–18, August 2021.
- [23] T. Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer, "Avalanche: A novel metastable consensus protocol family for cryptocurrencies," <https://www.avalabs.org/whitepapers>, 2020, accessed: 2024-07-06.
- [24] G. Bigini, V. Freschi, and E. Lattanzi, "Blockchain in the iot: Architectures and implementation," *Future Internet*, vol. 12, no. 12, p. 208, 2020, submission received: 3 November 2020 / Revised: 20 November 2020 / Accepted: 23 November 2020 / Published: 25 November 2020. [Online]. Available: <https://www.mdpi.com/1999-5903/12/12/208>
- [25] J. Noh, S. Jeon, and S. Cho, "Distributed blockchain-based message authentication scheme for connected vehicles," *Electronics*, vol. 9, no. 1, p. 74, 2020.
- [26] U. Hernandez-Jayo, A. S. K. Mammu, and I. De-la Iglesia, "Reliable communication in cooperative ad hoc networks," *Contemporary Issues in Wireless Communications*, 2014.
- [27] D. Manivannan, S. S. Moni, and S. Zeadally, "Secure authentication and privacy-preserving techniques in vehicular ad-hoc networks (vanets)," *Vehicular Communications*, vol. 25, p. 100247, 2020.
- [28] J. Meijers, P. Michalopoulos, S. Motepalli, G. Zhang, S. Zhang, A. Veneris, and H.-A. Jacobsen, "Blockchain for v2x: Applications and architectures," *IEEE Open Journal of Vehicular Technology*, vol. 3, pp. 193–209, 2022.

Stacking Ensemble Machine Learning Modelling for Milk Yield Prediction Based on Biological Characteristics and Feeding Strategies

Ruiming Xing^a, Baihua Li^{a*}, Shirin Dora^a, Michael Whittaker^b, and Janette Mathie^b

^aDepartment of Computer Science, Loughborough University, Loughborough, LE11 3TU, UK

^bCattle Information Service, Speir House, Stafford Park 1, Telford, Shropshire, WD3 3BD, UK

*Corresponding: b.li@lboro.ac.uk; {r.xing, s.dora}@lboro.ac.uk; {michaelwhittaker, janettemathie}@thecis.co.uk

Abstract—Knowing expected milk yield can help dairy farmers in better decision-making and management. The objective of this study was to build and compare predictive models to forecast daily milk yield over a long duration. A machine-learning pipeline was provided and five baseline models as well as a novel stacking model were developed for the prediction of milk yield on the CowNflow dataset using 414 Holstein cattle records collected from 1983 to 2019. Four different feature selection methods were performed to evaluate the essential biological characteristics and feeding-related features which affect milk yield. The results showed that the overall performance of predictive models improved after proper feature selection, with an R^2 value increased to 0.811, and a root mean squared error (RMSE) decreased to 3.627. The stacking model achieved the best performance with an R^2 value of 0.85, a mean absolute error (MAE) of 2.537 and an RMSE of 3.236. This research provides benchmark information for the prediction of milk yield on the CowNflow dataset and identifies useful factors such as dry matter (DM) intake and lactation month in long-term milk yield prediction.

Index Terms—Dairy Cattle, Milk Yield, Machine Learning, Feature Selection

I. INTRODUCTION

FORECASTING milk yield is a matter of great concern for dairy community. It has been shown that global milk demand is expected to grow by 22% between 2018 and 2027 [1]. It is important for dairy farmers to understand the essential factors influencing milk production so that they can deploy optimal management strategies, increase milk yield and reduce their production costs [2]. In this regard, predictive models for milk yield can help them develop better culling strategies and retain high-yielding cows [3].

In the past, livestock management relied more on the collective knowledge of people and their experience to make effective decisions. With the development of technology, dairy farmers are finding more effective management strategies such as using intelligent management systems, and sensors to record the characteristics of their herd and improve the efficiency of dairy production [4]. This has led to an increase in the availability of farm-related data, enabling data-driven management of farming through techniques like Machine Learning (ML).

Several studies have shown their interest in the field of milk prediction. Linear Regression (LR) [5], Multiple Linear

Regression (MLR) [6], Random Forest (RF) [7], [8], Support Vector Machine (SVM) [9] and Artificial Neural Network (ANN) [10] have been widely used in the prediction of milk yield. Sharma et al. [11] compared a multiple linear regression and ANN for milk yield prediction in Karan Fries dairy cattle and proved the performance of the ANN model is slightly superior to the regression model. While similar studies carried out in Sahiwal cattle [12] and Karan Fries cattle [13] also showed that ANNs gain good performance. Apart from basic factors such as cow age and lactation, Body weight at calving and the days in milk on the test day are regarded as the variables that are important for ANNs [11]. Other attempts to predict milk production involve a Back Propagation Neural Network (BPNN) optimised using Genetic Algorithm (GA) to analyse the impact of physiological and environmental, which was proposed by Sugiono et al. [14]. It is seen that predictive ML models have been deployed to deal with different scenarios. However, there is not much work being done to find the best-performing model from a machine-learning perspective to compare their performance in the same scenario.

Existing research on milk production has focused on accurately predicting milk yield over short periods. In [12], [13], neural networks are used to predict the milk yield from the first lactation 305-day. In [7], ANNs are deployed to predict milk yield for the first test day of the first lactation period. In addition, the XGBoost algorithm is applied for forecasting the next month's milk yield [15]. The above-mentioned works have not been evaluated for generating long-term predictions. A major reason for the lack of studies on long-term predictions is a lack of a suitable dataset.

Stacking is one of the most popular ensemble ML methods for predicting multiple nodes to build new models and improve model performance. It allows us to train multiple models to solve similar problems and build a new model with better performance based on their combined output [16]. In this study, a novel stacking method was proposed to accurately predict milk yield over a longer duration. Baseline models like LR, SVM, RF regression, AdaBoost and ANN were built and their performance was evaluated and compared. The

TABLE I
ATTRIBUTES RELATED TO COWS AND FEEDING CHARACTERISTICS

Index	Attributes	Num	Data Description	Data Type
1	Cow age (month)	414	Age in month	numerical
2	Body weight (kg)	414	Body weight	numerical
3	Physiological status	414	Two categories: dry, lactating	categorical
4	Lactation month	414	Number of months of lactation	numerical
5	Gestation month	414	Number of months of gestation	numerical
6	Diet type	414	Six categories, about feeding diet type	categorical
7	DM intake (kg/day)	414	Dry matter intake	numerical
8	DM digestibility (g/g)	414	Dry matter digestibility	numerical
9	DMI/100 kg body weight	414	Dry matter intake per 100 kg body weight	numerical
10	OM intake (kg/day)	413	Organic matter intake	numerical
11	Ash intake (kg/day)	413	Ash intake	numerical
12	N intake (g/day)	414	Nitrogen intake	numerical
13	CP intake (g/day)	414	Crude protein intake	numerical
14	Milk production (kg/day)	402	Milk production	numerical

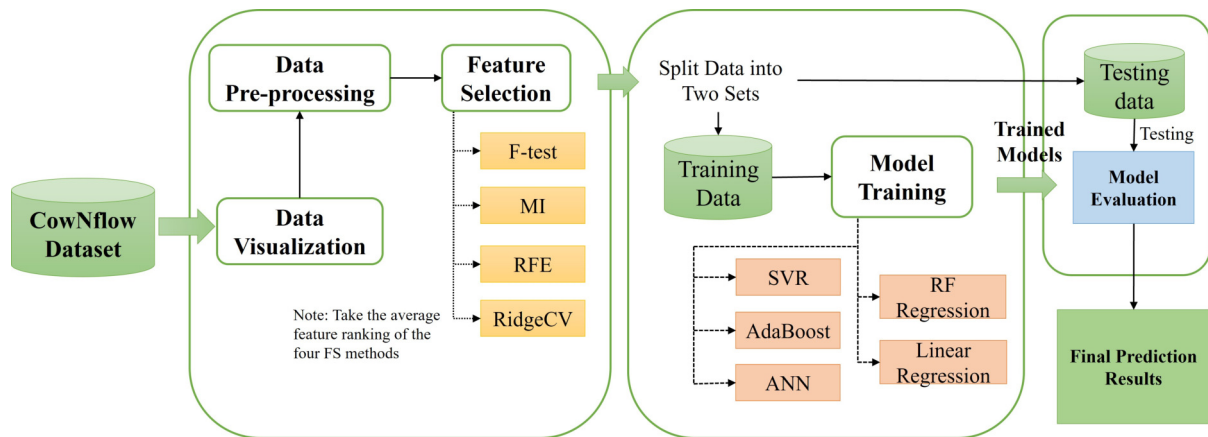


Fig. 1. ML pipeline for milk yield prediction

three main contributions in our study include: (1) Identify useful biological characteristics and feeding-related factors for long-term milk yield prediction. (2) Develop a stacking-based model for long-term milk yield prediction. (3) First build and compare the performance of various ML models on this dataset.

II. DATA DESCRIPTION

The data utilised in this study is called the CowNflow dataset [17] which is published by the National Institute for Agriculture, Food and the Environment (INRAE)¹ in France. The dataset has been collected at the experimental dairy farms of INRAE. It reports individual biological measurements from dairy cattle like dry matter (DM) intake, milk yield and feeding attributes like crude protein concentration of each feeding and diet composition. Cows were fed in individual troughs, had free access to water, and were milked twice a day. The dataset contains attributes like cow age, body weight, milk

yield, lactation number, feeding types and consumption of diet components. Table I shows the biological characteristics and feeding-related features in the dataset that are considered in this study.

III. EXPLORATORY DATA ANALYSIS

Figure 1 shows the pipeline for data analysis utilised in this paper. All the analysis reported in this paper has been carried out using the Python library Scikit-learn (version 1.1.3). Seaborn (version 3.10.6) is used for generating the visualisation.

Based on the dataset, data visualisation and pre-processing will be performed. After understanding the data distribution and cleaning the data, feature selection will be carried out. In feature selection, four different measures are taken into account. The ranking of the importance of features for milk prediction is obtained by averaging the ranking of features in each method. After that, the well-processed data is divided into training set and testing set. The training set is used to train the model and after getting the trained model and the

¹<https://entrepot.recherche.data.gouv.fr/dataverse/inrae>

testing set is used for model testing. Finally, the performance of each model will be evaluated and compared.

A. Data Visualisation

The purpose of visualisation is to develop an understanding of the underlying distribution for different features and identify patterns, and trends in the dataset.




Fig. 2. Histograms of part of cattle features (The vertical axis is the amount of records)




Fig. 3. Distribution of body weight vs. milk production

Figure 2 shows histograms for part of features such as cow age, body weight, and dry matter intake in the dataset. It can be seen that most of the cows in the experiment are between 2 and 10 years old. Histograms of body weight, nitrogen intake and milk production indicate that these features closely follow a Gaussian distribution, which may simplify the modelling process, and reduce the computational resources required for modelling.

The swarm chart (also named scatter plot) can be used to visualise the distribution of the joint distribution of a couple of discrete attributes. Figure 3 illustrates the impact of body weight on milk yield. It indicates that milk production increases as cow weight increases for cows weighing less than 600kg. For cows weighing more than 600kg, the milk yield doesn't exhibit a lot of variation.

B. Pre-processing

In this section, several data-cleaning steps are performed such as cleaning missing data and dealing with outliers to ensure the quality of data for further processes.

The records in which 'physiological status' has the value of 'dry' are not used in this study as these don't contribute towards predicting the milk yield. This resulted in 403 records that are used for further analysis in this paper.

Missing data: Dealing with missing data is important, as it may produce incorrect or biased results if missing data is not addressed properly. There are 3 missing values in each of the features *OM intake*, *Ash intake*, and *milk production*. These values are replaced by the mean values of the respective features.

Outliers: Many learning algorithms are sensitive to the range and distribution of attribute values. The interquartile range (IQR) is a commonly used tool to detect outliers with numerical values. To calculate the IQR, the dataset is divided into rank-ordered even quartiles, denoted by Q1 (lower 25%), Q2 (median 50%) and Q3 (upper 75% quartile), so IQR is the median 50% (Q3 – Q1). The whiskers have an offset length of 1.5*IQR, any data located outside of the whiskers is considered an outlier.




Fig. 4. Outlier detection using IQR box-and-whisker

Figure 4 shows the example box plots for different features in the dataset. The values outside of the whiskers are considered outliers. For example, For feature *DM intake*, the figure illustrates that values over 27 are regarded as outliers. In our task, to prevent loss of data available for training, only those records in which two or more features are identified as outliers are removed.

TABLE II
FEATURE RANKING WITH DIFFERENT FEATURE SELECTION METHODS

Features	F-test	MI	RFE	RidgeCV	Voting Rank
OM intake (kg/day)	1	1	2	2	1.50
DM intake (kg/day)	2	2	4	3	2.75
Lactation month	3	3	1	1	2.00
Diet type	4	7	11	4	6.50
DMI/100 kg body weight	5	8	5	8	6.50
CP intake (g/day)	6	5	9	10	7.50
N intake (g/day)	7	6	8	9	7.50
Gestation month	8	9	6	7	7.50
DM digestibility (g/g)	9	11	3	5	7.00
Body weight (kg)	10	10	12	6	9.50
Cow age (year)	11	4	7	11	8.25
Ash intake(kg/day)	12	12	10	12	11.50

Ranking score:1-12, 1 means most related and 12 represents the least.

C. Feature selection

Since irrelevant, redundant variables can reduce the model's generalisation capability and accuracy, feature selection is an effective step to find the most informative feature set that can have a better impact on the model performance [18]. Before selection. The 'diet type' feature was converted to a numerical feature using CatBoost Encoder [19].




Fig. 5. Heatmap on feature correlation in Pearson coefficients

A correlation heatmap shows the correlation coefficients between a set of variables. It can be especially useful to identify which variables are most strongly correlated with each other and define potential confounding factors. We convert the numbers to absolute values since correlation is shown by numerical magnitude, whereas positive or negative values simply indicate a positive or negative correlation. Then the correlation matrix with Pearson coefficients is shown in Fig. 5.

Four different measurements are applied to rank the features, which are linear regression f-test, Mutual information test, Recursive feature elimination (RFE) and Ridge regression with built-in cross-validation (RidgeCV). **F-test** is a statistical test which provides an f-score by calculating the ratio of variances. The variance of a feature determines how much it impacts the milk yield. If the variance is low, it implies this feature has less importance in predicting milk yield and vice-versa. **Mutual information** (MI) evaluates the gain of each variable in the context of the target variable, and it is predicated on joint probability. It indicates that the higher the mutual information value, the closer the connection between this feature and the target.

In addition to the two filter feature selected methods mentioned previously, a wrapper method **RFE** and an embedded method **RidgeCV** are also adapted to evaluate the correlation. RFE selects features by recursively considering smaller and smaller sets of features, the SVM algorithm and linear kernel were chosen to perform. RidgeCV is normally used in datasets which have multicollinearity. It uses L2 regularisation but performs Leave-One-Out Cross-Validation.

Table II shows the ranks of all the features based on the different methods used for feature selection. It can be observed that Ash-intake has the lowest average rank. Further, there is a difference of 3.25 between the rank of ash intake and the second lowest rank which indicates that ash intake was consistently ranked lower by all the feature selection methods. It can also be seen from Table II that CP intake and N intake have the same average rank. Further, it can be observed from Figure 5 that CP intake and N intake are strongly correlated with a Pearson coefficient of 1. Similarly, OM intake and DM intake are highly correlated with a Pearson coefficient of 1. Based on these observations, the features of ash intake, CP intake and OM intake are not selected for further analysis.

TABLE III
 COMPARISON OF LEARNING ALGORITHMS ON THE ORIGINAL AND SELECTED FEATURES

Model	No Feature Selection			With Feature Selection		
	R^2	MAE	RMSE	R^2	MAE	RMSE
Linear Regression	0.776	3.031	3.952	0.776	3.020	3.947
SVR	0.813	2.733	3.613	0.812	2.772	3.619
RF regression	0.805	2.920	3.684	0.804	2.902	3.695
AdaBoost	0.813	2.860	3.601	0.820	2.838	3.542
ANN	0.827	2.670	3.471	0.841	2.602	3.330
Stacking	0.843	2.568	3.308	0.850	2.537	3.236
AVERAGE	0.813	2.797	3.605	0.817	2.779	3.562

D. Model Training

After pre-processing and feature extraction, 397 records with 9 features are retained for developing ML models. Records are standardised before performed. The performance of all models is evaluated using hold-out [20] validation framework with 75% and 25% data used for training and testing, respectively.

Five different supervised ML techniques are considered to develop models for predicting milk yield, namely Linear regression (LR), SVM, Random Forest (RF) regression, Adaptive Boosting (AdaBoost) and Artificial Neural Network (ANN). All models are fine-tuned and evaluated to decide the best model. The linear regression algorithm is a basic and relatively common method for generating predictions, which is well understood and can be trained very quickly. For SVM, the kernel is set to 'RBF'. RF regression ensembles multiple decision trees into its final decision. Different numbers of trees are tested to determine that 100 estimators for RF achieved the best performance. AdaBoost is also an ensemble learning algorithm, it aggregates a set of weak classifiers into a strong classifier. We finally set the number of estimators to 100 and the learning rate to 0.5. For the ANN model, an input layer, a dense layer with 100 ReLU-activated neurons, and an output layer with Adam as its optimiser make up the neural network. After feature selection, a stacking method is proposed after building the five baseline models. The weak learners are made up of three best-performed models and the meta learner is set to Ridge regression. The structure of this model is shown in Fig. 6.

IV. EXPERIMENTAL RESULTS AND EVALUATION

In this section, the performance of different models developed in this study is evaluated and compared. The metrics used to measure the performance of the models include R^2 , Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). The R^2 is less than 1 where values close to 1 indicate that the model captures nearly all of the variation in the outcome of the target. MAE calculates the difference between predictive value and actual value for each data sample and takes the average absolute value of all samples. RMSE is similar to MAE but represents the square root of the average of squared errors in

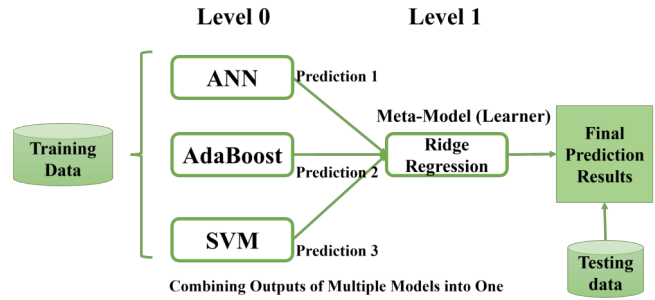


Fig. 6. The structure of stacking model

the predictions. The mathematical formulas for the different error metrics used in this study are given below:

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y}_i)^2}$$

$$MAE = \frac{1}{n} \sum |y_i - \hat{y}_i|$$

$$RMSE = \sqrt{\frac{1}{n} \sum |y_i - \hat{y}_i|^2}$$

where y_i represents the actual values, \hat{y}_i is the predicted values, \bar{y}_i is the mean of the actual data and n represents the total number of samples. A model that has smaller values of MAE and RMSE represents better performance.

Table III presents the performance of five baseline models on the dataset obtained before and after the feature selection. The table shows that when we remove the ash intake as well as the CP intake and OM intake, the overall performance of the trained models improves, with the average R^2 increasing from 0.807 to 0.811 and the values of both error measurements decreasing. It indicates that feature selection helps improve the performance of most models except SVM and RF regression. The R^2 of the best-performed model ANN increased from 0.827 to 0.840, and RMSE decreased from 0.441 to 0.426. The highest R^2 value obtained by the artificial neural network model was 0.827.

The three best-performed models are used to develop a model using the stacking technique. The performance of the

Stacking model is compared with the ANN model shown in Table III. Compared with the baseline models, the R^2 value of the stacking model improved to 0.85 and values of MAE and RMSE reduced to 2.537, and 3.236 respectively.

V. CONCLUSION

Milk production has received much attention in dairy farming. In this experiment, an ML pipeline is developed and applied to the CowNflow dataset for predicting milk yield. Four different feature selection methods were performed. 9 features of the original 13 were selected after data preprocessing and feature selection. Five different ML algorithms and a stacking method were utilized. Among the five baseline models, ANN achieved the best performance with a top R^2 value of 0.827 and the lowest RMSE of 3.471 before feature selection. After feature selection, the average values of R^2 for 5 models increased from 0.807 to 0.811, with both error measure matrices reduced. The stacking model had the best performance with an R^2 value of 0.85 and an RMSE value of 3.236.

According to the result, it is indicated that the *Ash intake* doesn't contribute much to the milk yield in long-term prediction. For the feeding factors, CP and OM intake are highly correlated to N and DM intake, respectively, which can be dismissed. The ML pipeline proposed in this study is shown to be efficient and generate good results. In future work, it can be optimised for further analysis and the current results will be a useful benchmark for further model comparison on this dataset.

VI. ACKNOWLEDGEMENT

The project is sponsored by Cattle Information Service (CIS) and National Bovine Data Centre (NBDC) in the UK. We also would like to thank UK Research and Innovation (UKRI) and the Engineering and Physical Sciences Research Council (EPSRC) for the support of this work (grant EP/Y00597X/1).

REFERENCES

- [1] OECD, Food, and A. O. of the United Nations, *OECD-FAO Agricultural Outlook 2018-2027*. OECD, 2018.
- [2] M. Cockburn, "Review: Application and prospective discussion of machine learning for the management of dairy farms," *Animals*, vol. 10, no. 9, 2020. doi: 10.3390/ani10091690
- [3] M. Lopez-Suarez, E. Armengol, S. Calsamiglia, and L. Castillejos, "Using decision trees to extract patterns for dairy culling management," in *Artificial Intelligence Applications and Innovations*. Springer International Publishing, 2018. doi: 10.1007/978-3-319-92007-8_20
- [4] S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, "Big data in smart farming—a review," *Agricultural systems*, 2017. doi: 10.1016/j.agsy.2017.01.023
- [5] A. Saha and S. Bhattacharyya, "Artificial insemination for milk production in india: A statistical insight," *Indian Journal of Animal Sciences*, vol. 90, no. 8, 2020. doi: 10.56093/ijans.v90i8.109314
- [6] F. Zhang, J. Upton, L. Shalloo, P. Shine, and M. D. Murphy, "Effect of introducing weather parameters on the accuracy of milk production forecast models," *Information Processing in Agriculture*, vol. 7, no. 1, pp. 120–138, 2020. doi: 10.1016/j.inpa.2019.04.004
- [7] G. M. Dallago, D. M. de Figueiredo, P. C. de Resende Andrade, R. A. dos Santos, R. Lacroix, D. E. Santschi, and D. M. Lefebvre, "Predicting first test day milk yield of dairy heifers," *Computers and Electronics in Agriculture*, vol. 166, p. 105032, 2019. doi: 10.1016/j.compag.2019.105032
- [8] M. Salamone, I. Adriaens, A. Vervaeke, G. Opsomer, H. Atashi, V. Fievez, B. Aernouts, and M. Hostens, "Prediction of first test day milk yield using historical records in dairy cows," *animal*, vol. 16, no. 11, p. 100658, 2022. doi: 10.1016/j.animal.2022.100658
- [9] Q. T. Nguyen, R. Fouchereau, E. Frenod, C. Gerard, and V. Sincholle, "Comparison of forecast models of production of dairy cows combining animal and diet parameters," *Computers and Electronics in Agriculture*, vol. 170, p. 105258, 2020. doi: 10.1016/j.compag.2020.105258
- [10] H. Radwan, H. El Qalioubi, and E. A. Elfadl, "Classification and prediction of milk yield level for holstein friesian cattle using parametric and non-parametric statistical classification models," *Journal of Advanced Veterinary and Animal Research*, vol. 7, no. 3, 2020. doi: 10.5455/javar.2020.g438
- [11] A. K. Sharma, R. Sharma, and H. Kasana, "Prediction of first lactation 305-day milk yield in karan fries dairy cattle using ann modeling," *Applied Soft Computing*, vol. 7, no. 3, 2007. doi: 10.1016/j.asoc.2006.07.002
- [12] V. Dongre, R. Gandhi, A. Singh, and A. Ruhil, "Comparative efficiency of artificial neural networks and multiple linear regression analysis for prediction of first lactation 305-day milk yield in sahiwal cattle," *Livestock Science*, vol. 147, no. 1-3, 2012. doi: 10.1016/j.livsci.2012.04.002
- [13] D. Njubi, J. Wakhungu, and M. Badamana, "Use of test-day records to predict first lactation 305-day milk yield using artificial neural network in kenyan holstein-friesian dairy cows," *Tropical animal health and production*, vol. 42, 2010. doi: 10.1007/s11250-009-9468-7
- [14] S. Sugiono, R. Soenoko, and D. P. Andriani, "Analysis the relationship of physiological, environmental, and cow milk productivity using ai," in *2016 International Conference on Data and Software Engineering (ICoDSE)*. IEEE, 2016. doi: 10.1109/ICoDSE.2016.7936165 pp. 1–6.
- [15] B. Ji, T. Banhazi, C. J. Phillips, C. Wang, and B. Li, "A machine learning framework to predict the next month's daily milk yield, milk composition and milking frequency for cows in a robotic dairy farm," *biosystems engineering*, vol. 216, pp. 186–197, 2022. doi: 10.1016/j.biosystemseng.2022.02.013
- [16] S. Džeroski and B. Ženko, "Is combining classifiers with stacking better than selecting the best one?" *Machine learning*, 2004. doi: 10.1023/B:MACH.0000015881.36452.6e
- [17] M. Ferreira, R. Delagarde, and N. Edouard, "Cownflow: A dataset on nitrogen flows and balances in dairy cows fed maize forage or herbage-based diets," *Data in Brief*, 2021. doi: 10.1016/j.dib.2021.107393
- [18] J. Li, K. Cheng, S. Wang, F. Morstatter, R. P. Trevino, J. Tang, and H. Liu, "Feature selection: A data perspective," *ACM computing surveys (CSUR)*, vol. 50, no. 6, 2017. doi: 10.1145/3136625
- [19] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "Catboost: unbiased boosting with categorical features," *Advances in neural information processing systems*, vol. 31, 2018. doi: 10.48550/arXiv.1706.09516
- [20] S. Raschka, "Model evaluation, model selection, and algorithm selection in machine learning," 2018. doi: 10.48550/arXiv.1811.12808

A bottom-up approach to select constrained spectral bands discriminating vine diseases

Shurong Zhang*, Alban Goupil*, Valeriu Vrabie*, Eric Perrin*, and Marie-Laure Panon†

*Université de Reims Champagne-Ardenne, CReSTIC, Reims, France

†Comité Champagne, 5 Rue Henri Martin, 51200 Épernay, France

shurong.zhang@univ-reims.fr; marie-laure.panon@civc.fr

Abstract—The detection and control of diseases constitute a primary objective of French viticultural research. In this paper, we present a bottom-up hierarchical approach for selecting spectral bands suitable for class discrimination of spectra acquired by Infrared spectroscopy. Our method entails evaluating neighboring bands using various similarity metrics, applying aggregation criteria, and ultimately identifying a limited number of the most relevant bands for the separation of classes. The bandwidths are limited within a range as is typically required for choosing existing optical filters or specifying colored filter arrays. Our approach facilitates the discovery of distinctive spectral bands associated with a disease of interest, enabling the customization of multispectral cameras to meet specific requirements. It was applied to spectra collected on vine leaves spanning a three-year period with the goal to identify the most discriminant bands for the detection of grapevine yellows. The results show that a limited number of bands are sufficient to identify this class of interest through a classifier based on Linear Discriminant Analysis.

Index Terms—Band selection, hierarchical model, classification, Grapevine Flavescence Dorée

I. INTRODUCTION

FLAVESCENCE Dorée, a serious epidemic disease, is one of the yellow diseases of grapevines. In order to control the risk of its spread, it is necessary to develop an effective and high-throughput detection tool. In this context, we are interested in developing an approach for selecting discriminative spectral bands based on spectra acquired by Near Infrared (NIR) spectroscopy, with a view to specifying a multispectral camera suitable for large-scale acquisitions. These spectral bands must allow optimal separation of the different classes, two in our application: yellow plants and healthy plants. On the other hand, we want to identify a limited number of spectral bands that can have different widths, but with restricted minimum and maximum bandwidths.

We position ourselves in the context of reducing the dimensions of spectral data. This goal can be achieved through band selection techniques that allow the selection of representative bands from an original spectrum. Although there are several selection methods, we will focus on those that are most relevant to our topic.

In [1], unsupervised hierarchical clustering approaches that merge similar wavelengths into bands were proposed. The models named WaluMi for mutual information and WaluDi

for Kullback-Leibler divergence use these metrics to assess the similarity between bands. The wavelength representing a band with the highest average similarity to others is selected as the band's representative. These models offer the advantage of merging wavelength similarities into a single band and choosing the optimal wavelength as the representative of this band. The main disadvantage is that only wavelengths are selected, which means that if one wants to specify an optical filter for a multispectral camera, it must have a very narrow band, which is very difficult to achieve in practice. Our approach focuses on aggregating multiple wavelengths into bands, thereby reducing computational complexity while maintaining the performance of identified bands. The rationale behind this choice is rooted in the physical correlation of adjacent wavelengths of NIR spectra, making the computation of similarities restricted to neighboring wavelengths sufficient. Additionally, the bands in our context are inherently adjacent wavelengths with minimal bandwidth. This deliberate restriction not only simplifies the computational requirements but also aligns with the practical constraints of our intended application.

A similar approach was proposed in [2], but using an adaptive hyperbolic distance as similarity measure. This distance avoids obtaining bands with a single or very limited number of wavelengths. Thereby, it partially respects our constraints, however, does not impose a maximum width of the bands. This is problematic in practice, particularly if we use the same sensor on which a colored filter array is placed, the quantity of absorbed light by the sensor being dependent on the bandwidth of each filter. Furthermore, this approach only calculates similarities restricted to neighboring bands.

In [3], the MRMR (Max Relevance and Min Redundancy) algorithm uses statistical selection to independently evaluate the importance of features. Its objective is to select an optimal subset that maximizes the relevance between features and classes, while minimizing the correlation between the selected features, achieved through the application of mutual information. This algorithm is typically applied to identify a set of discriminating wavelengths. In our case, we used this algorithm to identify discriminating bands after several wavelengths have been merged into bands.

In [4] an approach is proposed that uses a Random Forest method based on Recursive Feature Elimination (RFE) criteria. This approach is able to automatically eliminate feature redundancy and generally provides better and more compact subsets.

The project leading to this publication has received funding from the Grand-Est region, France.

The authors used Random Forest to evaluate the importance of features, and we have chosen it to compare it with MRMR in the second step of our method, where bands are selected. The feature selection of the RFE approach is also driven by the classes, but by a classification approach that may be different from the information theory used in MRMR. Note that there are similar approaches to RFE, such as Sequential Feature Selection [5], which greedily incorporates features by either adding or removing them to create a subset. Since fairly similar results were obtained with a Random Forest as an estimator, we only subsequently chose RFE for the sake of simplifying the display of the results.

These methods have been proposed to select wavelengths. Multispectral cameras typically capture a range of wavelengths with a central value and a certain number of wavelengths on either side. Those composed of several sensors use an optical filter for each sensor having fixed bandwidths imposed by the filter manufacturers, typically at least 10 nm. For specifying such a camera, it is possible to integrate the available NIR spectra with windows of desired width, such as Hanning, and then apply one of the previous methods by identifying bands which correspond to these integrations. However, all the filters will have the same bandwidth. For multispectral cameras using a single sensor with a color filter array, it is possible to specify the bandwidths of each filter but which, for practical reasons, must have a specified minimum and maximum limit.

The objective here is to propose an approach allowing aggregating similar wavelengths into bands and then identify a limited number of the most relevant bands for the separation of classes. In the first step, our hierarchical model allows multiple wavelengths to be combined into a band, emulating the physics of optical filters that can have different bandwidths, as may be the case for specifying color filter arrays. In the second step, a few bands are selected, as is necessary when one wishes to specify a suitable multispectral camera. Indeed, for cameras composed of several sensors, increasing the number of bands requires as many sensors and suitable electronics, which leads to cameras with important dimensions and weight and to a significant parallax phenomenon. For cameras using colored filter arrays, increasing the number of bands leads to a decrease in spatial resolution. This work is an extension of our previous work [6] in which we take into account new constraints, minimal and maximal bandwidths. The application is the same: specify a camera adapted to the detection of grapevine yellows by identifying optimal bands capable of effectively distinguishing healthy plants and grapevine yellows.

Since our method consists of two crucial steps: band grouping and band selection, we chose to use the spectral clustering algorithm studied in [7] to compare the performance of the band grouping. Spectral clustering is widely used to group data points based on their similarity. The author explores the concept of representing data as a graph, where connections between points reflect their similarity, and aims to divide the data into clusters such that points within a cluster are similar, while points in different clusters are dissimilar. In our approach, we adapt spectral clustering to group similar adjacent bands

by generating an adjacency matrix in which only adjacent bands are connected. This adjacency matrix is then passed to the spectral clustering algorithm with the affinity set to 'pre-calculated' to use the user-defined adjacency matrix. Once we have the groups of bands, we apply the same selection criteria as our approach to selecting the discriminative bands for class separation.

II. PROPOSED METHOD

As shown in Figure 1, the proposed hierarchical clustering approach is a two-step procedure. First, the wavelengths are merged based on their similarity by a bottom-up clustering into several bands that respect a width constraint, as presented in Section II-A. Each band is characterized by an interval that combines several weighted wavelengths to emulate the response of an optical filter. The pseudocode outline is shown in Algorithm 1. Following that, we assess the relevancy of the obtained bands for classes differentiation, as elaborated in Section II-B. A selection among these bands is done that is relevant according to the available classes. Finally, we choose the most effective bands for class discrimination, which are the final output of our algorithm, bands that allow the best classification, indicating whether they are relevant to the classes we have.

The inputs are the P NIR spectra, each acquired for a class $c \in C$ and having n wavelengths. A spectrum $S(\lambda_i) \in \mathbb{R}^P$ gives the reflectance of the target for each wavelength λ_i in the range $[\lambda_1, \dots, \lambda_n]$.




Fig. 1: The pipeline of our band selection approach.

A. Searching for relevant similarity while preserving adjacency

Initially, each band B_i consists of a single wavelength λ_i . The merging process is based on the similarity between adjacent bands, which is calculated on the representatives R_i of each band B_i . It begins by combining the most similar adjacent bands at each step and continues until no feasible band combination exceeds the specified maximum number Max of wavelengths. Here, Max represents the largest size of the desired optical filter response. If the merging of the most similar adjacent bands is constrained by Max , the algorithm will not merge them and consider them as independent bands, it will merge the next most similar adjacent bands that are not constrained by Max , and so on. This program accomplishes the task by calculating the distances between each pair of adjacent bands and generating a sorted list of bands. The algorithm attempts to merge bands, starting with the top selection and working down the list. As soon as two bands are

merged, the sorted list loop is closed, and the next iteration is started. At each iteration, it checks the total number of bands. If the number remains constant for two consecutive iterations, it indicates that the model cannot merge two bands without violating constraints. At this point the main loop is closed and the optimal set of bands is obtained. Additionally, each band should consist of at least Min wavelengths. At the end of the process, a check is performed and any bands that do not meet this criterion are removed.

Since similarity calculations are exclusively performed between a band and its two neighbors, the search for maximum similarity follows a linear pattern. To speed up the algorithm, similarities are not calculated at each iteration. Instead, the list of similarities between adjacent bands is computed at the beginning and is only updated between the new merged band and its neighbors in each iteration. Compared to the bottom-up hierarchical groupings mentioned above, this implementation is more adaptable to our case for selecting the groups of neighboring wavelengths. It also has quadratic complexity compared to cubic complexity.

Data: List of wavelengths $\lambda_1, \dots, \lambda_n$
Data: Reflectance $S(\lambda_i) \in \mathbb{R}^P$, with P the number total of spectra, each belonging to a class $c \in C$
Result: Set of m bands $\mathcal{B} = \{B_1, \dots, B_m\}$
 /* Initialization */
 $B_i = \{\lambda_i\}, i = 1, \dots, n;$
 Representatives $R_i = S(B_i)$ cf section II-A1;
 Distances $d(R_i, R_{i+1})$ between the representatives of bands adjacent, $i = 1, \dots, n - 1$ cf section II-A2;
 /* Hierarchical grouping */
while $length(B) > 1$ **do**
 $list_sort \leftarrow argsort_i d(R_i, R_{i+1});$
 $bandsize = length(B);$
 while $a \in list_sort$ **do**
 if $length(B(a)) + length(B(a + 1)) < Max$
 then
 $B_a \leftarrow B_a \cup B_{a+1}$ and remove B_{a+1} from $\mathcal{B};$
 Compute the representative $R_a = R(B_a)$ of $B_a;$
 Update the distances $d(R_{a-1}, R_a)$ and $d(R_a, R_{a+1});$
 /* adjacent bands found */
 break;
 if $length(B) = bandsize$ **then**
 /* no aggregation found */
 break;
for i **from** 1 **to** $length(B)$ **do**
 if $length(B[i]) < Min$ **then**
 Delete $B[i];$
return $\mathcal{B};$
Algorithm 1: Our hierarchical bottom-up classification

1) *Representative of a band:* Our approach consists of weighting spectral information at wavelengths composing a band B_i . We use a Hanning window that emulates a real optical filter physically, but other windows can also be used. The R_i representative of the B_i band is calculated as the weighted average of the spectral information in this window

$$R_i = \frac{1}{L} S(B_i) \cdot W(L), \quad (1)$$

where $W(L) \in \mathbb{R}^L$ is the Hanning window of size L , corresponding to the number of wavelengths contained in the band B_i , and $S(B_i) \in \mathbb{R}^{P \times L}$ is the spectral information at the wavelengths contained in this band.

2) *Distance between the representatives:* We estimate the distance between the representatives of these bands as the similarity between these bands. Four distances between band representatives are considered in our work: the commonly used Euclidean distance D_{L2} ; the adaptive hyperbolic distance D_{HY} proposed by [2]; the Jensen-Shannon divergence D_{JS} , which takes advantage of the distribution of spectral information rather than spectral values; and the conditional mutual information D_{CMI} , which takes into account the labels $c \in C$ associated to each spectrum. Only the last two are described in detail below.

a) *Jensen-Shannon (JS) divergence:* can be interpreted as a kind of similarity between two probability distributions that symmetries the Kullback-Leibler distance. We assume that R_i and R_j are the representatives of the bands B_i and B_j . The distance based on the divergence of JS, D_{JS} is defined by

$$D_{JS}(R_i, R_j) = \frac{1}{2} D_{KL}(R_i | M) + \frac{1}{2} D_{KL}(R_j | M), \quad (2)$$

where D_{KL} is the Kullback-Leibler divergence between the probability densities of the variables R_i, R_j and M , the mean distribution of the distributions of R_i and R_j . The Jensen-Shannon divergence is always positive or zero. It cancels out when R_i and R_j have the same probability distribution. Thus, two bands B_i and B_j with a low value of D_{JS} are combined if the representatives of these bands have close densities.

b) *Conditional Mutual Information (CMI):* measures the dependence between the spectral information of the B_i bands in each class $c \in C$. Each class c has a probability $p(c)$ corresponding to its proportion among the data set. The mutual information $I(R_i; R_j | c)$ between the representatives of B_i and B_j conditioned to class c is given by the weighted average

$$I(R_i; R_j | C) = \sum_{c \in C} I(R_i; R_j | c) \times p(c). \quad (3)$$

and the CMI distance is defined by

$$D_{CMI}(R_i, R_j) = \frac{1}{1 + I(R_i; R_j | C)}. \quad (4)$$

This distance offers a key advantage because it incorporates spectral information from two bands within different classes.

B. Selecting relevant bands for classification

Since the bandwidth of the built bands is constrained, the procedure may yield too many representative bands.

1) *Examining the variance (VAR)*: Examining the variance between each band and the classes allows us to prioritize bands with higher values. The ranking of them involves computing the variance between the final representative of each band and the class, and the selection of discriminating representatives is based on this ranking.

2) *Leverage state-of-the-art methods*: Using band selection techniques such as MRMR and RFE, we are able to select discriminating representatives with respect to the class.

We chose these selection criteria because they allow us to evaluate the variance, mutual information and accuracy of the classifier, three different aspects that allow us to evaluate the discriminatory power of the bands identified for class separation. Once we have found those bands, given that they are the most relevant bands for class separation that we have, we can then specify the filters of a multispectral camera suitable for the application considered.

III. RESULTS

We have carried out several acquisition campaigns on vine leaves of the Chardonnay grape variety at the Comité Champagne's Plumecoq experimental estate. The acquisitions were carried out after the harvest periods of the years 2021 to 2023. NIR spectra were acquired with a LabSpec4 portable spectrometer on around 700 grapevine leaves per year that were picked and arranged on polystyrene boards. Two spectra were acquired on each leaf, which gives a collection of $P = 4282$ spectra acquired between 2021 and 2023. This spectrometer provides spectral information every 1 nm from 350 to 2500 nm, but only the 400 – 1000 nm region was retained. The collection $S(\lambda_i) \in \mathbb{R}^P$ thus represents the reflectance of the leaves for each wavelength λ_i in the range $[\lambda_1, \dots, \lambda_n]$ with $n = 600$. The aim of this study was to identify the spectral bands that allow us to distinguish the $C = 2$ classes: healthy plants and grapevine yellows. To identify the discriminating bands, from the collection of $P = 4282$ spectra, 1709 spectra correspond to grapevine yellows and 2573 spectra to healthy plants.

We set $Min = 10$ wavelengths and $Max = 100$ wavelengths because the bands of the optical filters available on the market or the color filter array that can be specified are generally within these limits. All available $P = 4282$ spectra were used to identify discriminative bands with the proposed approach. For comparison, as mentioned in the introduction, we adapted the spectral clustering to group the adjacent wavelengths into bands, without bandwidth restriction, and then applied the selection techniques to obtain the discriminative bands.

Once the bands have been identified, we divided the data, i.e. the spectral information at the selected bands, into two sets, one training and one test. The training set contains 90% of the data, i.e. 3854 spectra, and the test set contains the remaining 10%, i.e. 428 spectra. In the former set, a Linear Discriminant Analysis (LDA) classifier was trained, and classification performance was evaluated in the latest set considering spectral information at the bands identified by

the selection methods. We decided to use the LDA classifier for the evaluation because, compared to other classifiers such as Support Vector Machine (SVM) or k-Nearest Neighbors (KNN), LDA is more often used for feature extraction, where it identifies the most relevant features for class discrimination.




Fig. 2: Band grouping and VAR selection.




Fig. 3: Band grouping and MRMR selection.

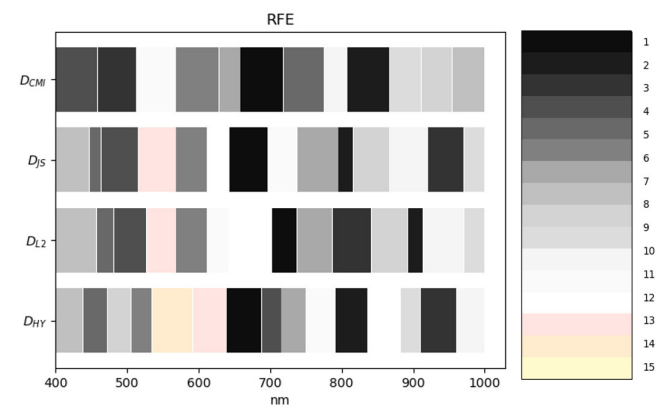


Fig. 4: Band grouping and RFE selection.




Fig. 5: Accuracy values for LDA with all methods.

	D_{HY}	D_{L2}	D_{JS}	D_{CMI}
VAR	0.8659	0.8631	0.8785	0.8757
MRMR	0.8617	0.8778	0.8794	0.8619
RFE	0.8799	0.8848	0.8834	0.8808

(a) 6 discriminating bands

	D_{HY}	D_{L2}	D_{JS}	D_{CMI}
VAR	0.8846	0.8787	0.8790	0.8757
MRMR	0.8610	0.8778	0.8820	0.8764
RFE	0.8794	0.8893	0.8846	0.8811

(b) 7 discriminating bands

Fig. 6: Accuracy of LDA for 6 and 7 discriminating bands

Figure 2, 3 and 4 show the bands obtained in relation to the different distances proposed in section II-A2. The colors assigned convey the discriminatory quality of the bands for class discrimination in relation to the different criteria proposed in section II-B. A darker color indicates the discriminatory nature of a band and the accompanying number indicates its

position as the n^{th} band selected by the specified method. A consistent trend is observed across all methods, showing a higher frequency of selection for bands around 700 nm and an intermediate frequency of selection for bands between 400 nm and 500 nm across all three selection methods. In contrast, bands between 900 nm and 1000 nm are rarely selected in the MRMR selection. We could also note that, as discussed in the introduction, D_{HY} has more uniform bands than the other methods. In addition, each distance provides a different number of bands: D_{HY} has 15 bands, D_{L2} and D_{JS} each have 14 bands, and D_{CMI} has 12 bands.

In Figure 5, the LDA accuracy is shown for each method, with distinctive markers assigned based on the selection technique: variance (VAR) selection is marked with 'o', MRMR selection is marked with '△', and RFE selection is marked with '*'. Different colors are also assigned to each distance metric: red for the hyperbolic distance (D_{HY}), green for the L2 distance (D_{L2}), purple for Jensen-Shannon distance (D_{JS}), and cyan for conditional mutual information distance (D_{CMI}). Straight curves are used for spectral clustering with

selection techniques. The results show that increasing the number of bands is associated with an expected increase in LDA accuracy due to the additional spectral information. Note that the accuracy of LDA using all 600 available wavelengths (from 400 to 1000 nm) is 0.9523. In particular, techniques that use D_{JS} , D_{L2} , and D_{CMI} show superior performance over different numbers of bands. In contrast, methods that use D_{HY} show less effectiveness, especially when we select a small number of bands, possibly due to the narrow discriminating bands at other distances compared to the uniform nature of all bands in D_{HY} .

However, the hyperbolic distance works well in our previous methods [6], which is to be expected because, as we mentioned in the introduction, this distance produces almost uniform bands. However, even though it is designed to avoid this type of problem without imposing a bandwidth limit, this distance leads to quite heterogeneous bands in terms of width. Even if the classification results obtained with LDA are satisfactory, the practical realization of a multispectral camera using a color-array filter could still be difficult, since the amount of light absorbed by the same sensor depends on the bandwidth of the filters.

For spectral clustering, we regrouped the adjacent wavelengths in 15 bands using the spectral clustering, then we applied the three criteria of selection. We could notice that its performance is less than our methods except for RFE-Spectral clustering in the case of selecting 5 discriminating bands.

Compared to the original hierarchical clustering method, our model imposes limits and stops as soon as we can't merge bands within the size limit. This approach prevents us from producing a band that exceeds the camera's range. However, spectral clustering does not have the constraint of the camera's band range, meaning it is possible to obtain a band that falls outside the camera's range in the results.

The evaluation is limited to 1 to 10 bands, in line with the objective of constructing multispectral cameras with relevant characteristics, including dimension, volume, and spatial resolution. Beyond this threshold, the associated cost increases significantly and widely available cameras typically offer between 6 and 7 band configurations. Consequently, figure 6 summarizes the accuracy for two classical configurations. Using 6 and 7 representative bands, the most effective method

is RFE- D_{L2} with an accuracy of 0.8848 and 0.8893.

IV. CONCLUSION AND OUTLOOK

We have proposed a new hierarchical bottom-up approach that merges the bands and then selects the most discriminating ones. This approach, which imposes a constraint on the bandwidths, aggregate similar bands by assessing the similarity between band representatives using different distances (L2, hyperbolic distance, Jensen-Shannon divergence, conditional mutual information). Finally, it selects the most relevant bands for the separation of classes using criteria based on maximization of the variance between bands and classes or using state-of-the-art selection techniques, MRMR and RFE.

This approach is useful in specifying a multispectral camera suitable for a specific application. It requires spectra to be recorded on samples belonging to different classes and to indicate the number of desired bands. In our case, we used it to identify bands distinguishing the vine yellows from healthy plants.

Work is underway to make the algorithm robust to multi-year acquisitions and to configure it so that it identifies bands that are insensitive to the acquisition year.

REFERENCES

- [1] A. Martínez-Usó, F. Pla, J. M. Sotoca, P. García-Sevilla, Clustering-Based Hyperspectral Band Selection Using Information Measures, *IEEE Trans. Geosci. Remote Sens.* 45(12), pp. 4158-4171, Dec. 2007. DOI: 10.1109/TGRS.2007.904951
- [2] H. Sun, L. Zhang, J. Ren, H. Huang, Novel hyperbolic clustering-based band hierarchy (HCBH) for effective unsupervised band selection of hyperspectral images, *Pattern Recognition*, 130 :108788, October 2022. DOI: 10.1016/j.patcog.2022.108788
- [3] H. Peng, F. Long, C. Ding, Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy, *IEEE Trans. Pattern Anal. Mach. Intell.* 27(8), pp. 1226-1238, Aug. 2005. DOI: 10.1109/TPAMI.2005.159
- [4] P. Granitto, C. Furlanello, F. Biasioli, F. Gasperi, Recursive feature elimination with random forest for PTR-MS analysis of agroindustrial products, *Chemometrics and Intelligent Laboratory Systems*, 83(2), 2006, DOI: 10.1016/j.chemolab.2006.01.007
- [5] D. Aha, R. Bankert, A comparative evaluation of sequential feature selection algorithms, *Proc. Fifth International Workshop on Artificial Intelligence and Statistics*, PMLR, 1995.
- [6] S. Zhang, E. Perrin, A. Goupil, V. Vrabie, M.-L. Panon, Nouveau modèle hiérarchique ascendant pour la sélection des bandes spectrales discriminant les maladies de la vigne, *Colloque GRETSI*, Grenoble, 2023.
- [7] U. V. Luxburg, A tutorial on spectral clustering, *Statistics and Computing*, 17(4):395-416, December 2007. DOI: 10.48550/arXiv.0711.0189

Dynamic Relationship Between Population Densities and Air Quality in the Four Largest Norwegian Cities

Petar Zhivkov

0000-0001-5687-5277

Inst. of I&C Tech., Bulgarian Academy of Sciences
Sofia, Bulgaria

Email: zhivkovpetar@gmail.com

Todor Kesarovski

0000-0002-0971-7783

Dep. of Safety, Economics & Planning
University of Stavanger, Norway

Email: toдор.m.kesarovski@uis.no

Abstract—Air pollution is a significant cause of health problems and disease worldwide. Considering the rapid urbanisation at a global scale in recent decades, resulting in more and more people in urban areas, cities deserve special attention in this regard. In this paper, we use air quality measurement data from 2010 to 2023 in the four largest Norwegian cities (Oslo, Bergen, Trondheim, and Stavanger) and correlate it with the evolution of population densities for the same period. The empirical analysis focuses on nitrogen dioxides (NO₂) and particulate matter (PM_{2.5} and PM₁₀) as critical pollutants in urban areas to verify whether their concentrations are affected by the increase in population densities for individual municipalities. In addition, we also correlate the data on air pollutants with different natural indicators such as temperature, air pressure, humidity, wind, and the rate of motorisation in the cities of interest.

I. INTRODUCTION

SINCE the rise of sustainability concerns in the 1970s, the focus on the negative impact of cities on the environment is growing gradually and air pollution is considered a major one. A comparative study between 20 cases in various European countries by [8] demonstrates that air pollution levels in Norway are similar to the other Scandinavian countries but lower than in southern Europe. The Norwegian Institute of Public Health (NIPH) confirms this statement in its annual health report from 2018 [18]. The institute documents that the air pollution levels in Norway have been relatively stable over the last decade, as the levels of key pollutants particulate matter (PM) and nitrogen dioxide (NO₂), have been even slightly declining.

However, it seems crucial to measure the extent of the negative anthropogenic factors of urban development and concentrations of activities on the accumulation of hazardous air pollutants within this discussion. Recent studies, including [11] and [3], emphasise this negative correlation, respectively, in the context of China and Germany. In Norway NIPH also outlines that the concentration of pollutants, such as PM and NO₂ varies considerably between urban areas and elsewhere in the country. There may also be significant variations within

The work is supported by the National Scientific Fund of Bulgaria under the grant DFNI KP-06-N52/5.

each town and densely populated area, depending on traffic and other emissions [18]. This is a fundamentally critical question considering the health, economic, and environmental impact that air pollution has, defining the latter as an essential determinant of the quality of life [3]. Examples of behaviours that contribute to air pollution include driving vehicles that emit pollutants, idling engines, burning fossil fuels for heating and energy, and engaging in activities that produce emissions like industrial processes and releasing harmful chemicals into the air [5].

On the other hand, it is crucial to recognise that denser urban environments can provide more effective use of resources such as space, energy, and raw materials. Denser cities require less space, saving valuable peri-urban lands instead of providing environmental services and agricultural goods. By reducing travel distances and enhancing connectivity through (a combination of) public transport, cycling, or walking, it is possible to reduce fossil fuel use significantly. More compact buildings also require less energy for cooling and heating [15]. Moreover, dense urban environments provide the conditions to reduce and improve supply chain coordination [13], leading to a better logistic distribution of goods [1]. Consequently, denser built environments can reduce CO₂ emissions, resulting in densification as a practical approach to guiding sustainable urban development [14].

Motivated by this somewhat twofold impact of urban densities, we address the relationship between population densities and concentration of air pollutants in the context of the four largest cities in Norway, i.e., Oslo, Bergen, Trondheim, and Stavanger. This goal defines the main research question of this study: What is the correlation between the population density's dynamics and the concentration of air pollutants within the four largest Norwegian cities?

By exploring this question, we aim to outline the correlation between the concentration of people, respectively, human activities and air quality. Therefore, we also make a modest attempt to compare the correlation between air pollution and population densities with the correlation between air pollution and other geographic-meteorological factors, such

as wind speed, temperature, humidity, atmospheric pressure, precipitation, and sun irradiation (sunshine hours).

II. AIR QUALITY MONITORING AND NORMS REGARDING SPECIFIC POLLUTANTS

Measurements of air emissions are critical for epidemiology and air quality control. Traditionally, concentrations of air emissions have been measured by air monitoring stations with standard equipment, allowing for highly reliable results. Although recently there is a scientific concern about the limitations of this, ground-based air pollution observations have limitations in terms of assessing personal exposure to pollution, this traditional approach possesses strong capacities in estimating the extent of air pollution on a city and municipal level [21], [12]. Thus, data collected by stationary monitoring stations in Norwegian cities fit the purpose of the study.

A. Hazardous pollutants of interest

Air pollution consists of a range of different substances, depending on the source. However, three of them, i.e., PM, NO₂ and ozone (O₃), are considered to be the most critical air-polluting components in urban areas that cause problems, disease, and death in this part of the world [4]. This study focuses on the former two – PM and NO₂. PM air pollution is a suspended combination of solid and liquid particles that vary in quantity, size, shape, surface area, chemical composition, solubility, and origin. Total Suspended Particles (TSPs) have a trimodal distribution in the ambient air, including coarse particles, fine particles, and ultrafine particles [20]. PM size-selective sampling refers to collecting particles below, above, or within a defined aerodynamic range of sizes, which is commonly chosen to be relevant to inhalation, deposition, causes, or toxicity [6]. Particle size is generally described in terms of a 50 per cent cut point at a specific aerodynamic diameter (such as 2.5 or 10 μm) and the slope of the sampling effectiveness curve because samplers are incapable of accurate size distinction. In terms of establishing guidelines or standards for acceptable levels of ambient PM pollution, public health policy has primarily focused on indicators of fine particles (PM_{2.5}), inhalable or thoracic particles (PM₁₀), and thoracic coarse particles (between PM₁₀ and PM_{2.5}) [9], [19]. Nitrogen Dioxide (NO₂) is another pollutant, part of a group of reactive gases known as oxides of nitrogen or nitrogen oxides (NO_x). Other nitrogen oxides include nitrous acid and nitric acid. NO₂ is used as the indicator for the larger group of nitrogen oxides [7]. In cities, NO₂ primarily accumulates in the air from internal combustion engines burning fossil fuels, i.e., motor vehicles, power plants, and offroad equipment [4]. Exposures to NO₂ over a short period can aggravate respiratory diseases, particularly asthma, leading to respiratory symptoms (such as coughing, wheezing or difficulty breathing), hospital admissions and visits to emergency rooms. Longer exposures to elevated concentrations of NO₂ may contribute to the development of asthma and potentially increase susceptibility to respiratory infections. People with

asthma and children and the elderly are generally at greater risk for the health effects of NO₂ [9], [19].

B. Norms regarding specific air pollutants

Although providing a solid background regarding limiting air pollution in urban and rural areas, WHO guidelines can be considered practical recommendations without mandatory character. On the other hand, in Europe, the EU Air Quality Directive represents legislation that every EU member must follow. The most noticeable difference is that there is no limit for daily limits for PM_{2.5} in the EU Air Quality Directive. In Norway, however, there are implemented limit values for each of the pollutants of interest, based on a long-term national ambition for local air quality. These goals are set to the same level for each city as air quality criteria for PM (annual mean) and NO₂ (annual mean). These criteria are established so that most of the population in the country is effectively protected against harmful effects if they are exceeded [18]. Based on newer knowledge about the health effects of PM by NIPH, the criteria for air quality concerning PM₁₀, PM_{2.5}, and NO₂ was revised and set to lower limits in 2014 [17], [16]. The current limits for the pollutants of interest in Norwegian cities are as follows: for PM₁₀ [max: 25 $\mu\text{g}/\text{m}^3$, target: 20 $\mu\text{g}/\text{m}^3$]; for PM_{2.5} [max: 15 $\mu\text{g}/\text{m}^3$, target: 8 $\mu\text{g}/\text{m}^3$]; for NO₂ [max, annual mean: 40 $\mu\text{g}/\text{m}^3$, max, hourly value: 200 $\mu\text{g}/\text{m}^3$].

III. METHODOLOGY

The method applied in this study is a quantitative approach to explore the relationships of interest regarding the concentration of air pollutants. This section describes the various aspects of the method, including the origin of the utilised datasets, their specifics and the approach through which correlations of interest are estimated.

A. Data

1) *Air quality data:* Regarding the emissions data of PM₁₀, PM_{2.5}, and NO₂, we use the annual mean values from 2010 to 2023. These datasets are extracted from the Central database for local air quality (SDB) and collected by the Norwegian Institute for Air Research (NILU). These stations are representative of NILU and stationary, as there are a different number of measurement stations located in the four cities. However, they are positioned in strategic locations to record air pollution in key areas, lying on streets and transport axes, or dispersed throughout cities. These variables are added to our regressions as dependent variables correlated with the weather, and density indicators.

2) *Population densities:* Population density (POPd) is a well-established indicator in urban planning, and its use is documented in the works of notable researchers, such as [22], [23], [10], among others. The measurement focuses on estimating the number of people per spatial unit. It is predominantly calculated in people per hectare and in "gross densities," as in the case of this study. The dynamics of the variables are calculated for the period 2010-2023 for the municipalities of interest, based on the census data publicly provided by the

Norwegian Statistics Bureau (SSB). The specific estimations are calculated through the following formula:

$$POPdx = POPx/Ax \quad (1)$$

where:

POP = total population

A = base land area (in hectares)

x = area of aggregation (the municipalities of Oslo, Bergen, Trondheim, and Stavanger)

3) *Weather data:* As argued in the previous section, the ambient concentrations of emissions tend to be affected crucially by the specific weather conditions in each context [2]. Thus, although not representing the primary research direction of interest for this study, we still use weather variables to generate basic linear regressions with each one of them. The measurements we include are the annual means of temperature (°C), atmospheric pressure (hPa), humidity levels, wind speed (m/s), average precipitation (mm), average precipitation days (≥ 1.0 mm), and average sunshine hours. For the first four indicators, we use the collected data for 2023 by all stations, which we use for extraction of emissions data, for the precipitation data, we use the aggregated data of NILU for the specific cities for the period of 1991-2023, and for the average sunshine hours, we employ the data for the period 2016-2023.

B. Correlation analysis

The last step of the data processing consists of performing statistical correlations between the concentration of air pollutants and population densities for the four case studies. This activity is realised through a series of correlation analyses (Pearson's correlation coefficient). The measure represents a practical statistical approach to exploring the relationship between two variables based on their values' standard deviation. This defines it as a normalised bivariate measurement whose value is always between -1 and 1 .

$$r = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \quad (2)$$

where:

r = correlation coefficient

x_i = values of the x-variable in a sample

\bar{x} = mean of the values of the x-variable

y_i = values of the y-variable in a sample

\bar{y} = mean of the values of the y-variable

Based on these estimations, we explore the relationships of interest by executing the correlations between the annual mean of each pollutant and the population density for the respective year for the period 2010-2023, excluding two correlations for which data is not available. This provides 42 correlations based on which basic regression analysis is elaborated. As a last verification activity to compare the influence of the population densities with the impact of weather indicators on air pollution, additional regressions are also executed. They correlate the extracted weather data (annual mean average) with the mean yearly average for the concentration of air pollutants (2010-2023) for the respective cities.

IV. RESULTS

The following section presents the generated outcomes from the undertaken data processing and discusses the empirical assumptions that can be drafted based on them.

A. Case studies

Before presenting the results from the data processing and the executed correlations, it is required to present a brief description of the examined cases. This is an essential element that allows us to interpret our results more rigorously. As a part of the case description, we include:

- The number and locations of the stationary sensors of NILU, based on which the data regarding air pollution is collected;
- The aggregated weather data for each city;
- The dynamic of the population density of its municipality for the period of 2010-2023.

1) *Oslo:* Oslo is the most carefully monitored case of the four examined as the capital and most populous city in Norway. NILU has 14 positioned stations all around Oslo's municipality, from where data is permanently collected. Oslo is also a specific case since it is the only one of the four major Norwegian cities not on its west coast. Thus, it is characterised as the case with the least average annual mean of humidity, precipitation, and wind but the city with the most sunshine hours annually. The municipality of Oslo is also the most densely populated one in Norway and, therefore, in between all of the four examined cases with a density of 54 people per hectare in 2023, see Table I. Since 2010 we can observe a gradual increase in the population density of approx. 2% each year. Exclusions are observed in 2011 and 2013 when there is a more significant increase and in 2017, 2021, and 2023 when there is no increase.

2) *Bergen:* Bergen is the second-most populous city in Norway. The emissions and weather data for Bergen are collected through five stations located all around the municipality. Regarding weather specifics, it is worth outlining that Bergen is by far the rainiest city of the examined cases. Its annual mean for average precipitation is double the one of Stavanger and triples the value of Oslo and Trondheim. This condition respectively reflects on the highest humidity and lowest sunshine hours from the examined cases. Although being the second-most populous city in Norway, the municipality of Bergen is characterised by lower population density numbers (with 30 people per hectare in 2023) even with respect to Stavanger and Trondheim. Since 2010 we can indicate a state of stagnation due to the stable number. Looking at the numbers, we can observe two rapid jumps of the density indicator in 2013 and 2020, but this is due to administrative changes resulting in a reduction in the municipality's area. Such restructuring events affect air pollution data records, as monitoring stations within added or removed municipal territories are included or excluded from the aggregated data.

TABLE I
POPULATION DENSITIES (INHABITANTS PER HECTARE) FOR THE FOUR LARGEST MUNICIPALITIES IN NORWAY, 2010-2023

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
OSLO POPd	42	44	45	47	48	49	50	50	51	52	53	53	54	54
OSLO +/-	-	+4.6%	+2.2%	+4.3%	+2.1%	+2.0%	+2.0%	-	+2.0%	+1.9%	+1.9%	-	+1.9%	-
BERGEN POPd	23	23	23	27	27	28	28	28	28	28	30	30	30	30
BERGEN +/-	-	-	-	+14.8%	-	+3.6%	-	-	-	-	+6.7%	-	-	-
TRONDHEIM POPd	25	25	26	30	30	30	30	31	31	31	32	33	33	34
TRONDHEIM +/-	-	-	+3.8%	+13.3%	-	-	-	-	-	-	+3.1%	+3.1%	-	+3.1%
STAVANGER POPd	26	27	27	31	30	31	31	31	31	31	17 (31)	17	17	17
STAVANGER +/-	-	+3.7%	-	+27.0%	-3.2%	+3.2%	-	-	-	-	-45.1% (-)	-	-	-

3) *Trondheim*: Trondheim is the most northern city of all examined cases, and its population is lower than Bergen's. NILU collects the climate data in the municipality via four permanent stations. Despite its geographic location to the north and on the coast, Trondheim has specific topographic conditions that result in lower humidity, wind, and precipitation compared to all other cases on the west coast. The municipality of Trondheim represents the second-most dense case from the ones examined in this study, with a density of 34 people per hectare for 2023. Looking holistically at the number, we can state that they illustrate a minor increase, as the only rigid one (in 2013) is due to administrative restructuring of the municipal borders.

4) *Stavanger*: Stavanger is a comparable city with Trondheim in terms of population size but weather-wise with Bergen. However, the degree of precipitation is still significantly lower than the one recorded in Bergen. Apart from this, it is worth mentioning that Stavanger is the warmest and the windiest city of all examined cases. In Stavanger, the weather and air pollution data are collected by four stations positioned within the municipality. By looking at the population densities dynamic we can argue that the case of Stavanger represents a state close to stagnation with fluctuating values. Similar to the case of Bergen, the two dramatic changes in the number (in 2013 and 2020) are due to administrative restructuring. However, the latter results in a significant decrease in the population density to 17 people per hectare. This is due to the incorporation of two low-dense populated municipalities within the administrative body of Stavanger.

B. Air quality

As section 2 presents in greater detail, there are different values to assess air quality levels. However, due to health effects and the Norwegian policies in this study, we focus on NO₂, PM_{2.5} and PM₁₀ as hazardous air pollutants. In this sub-section, the dynamics of their annual means (average annual concentrations from 2010 to 2023) are presented.

1) *Nitrogen dioxide (NO₂)*: Table II illustrates the dynamic development of the annual means throughout the period of interest for all cases. By examining the data, we can identify the apparent trend of reducing the emission of NO₂ in all of the examined cases in the last decade. This trend is even more evident after 2016 when the values for the instances of

Oslo and Bergen had been fluctuating. From a contemporary perspective, the levels for NO₂ do not exceed the standard max, namely, an annual mean of 40 $\mu\text{g}/\text{m}^3$.

2) *PM₁₀*: The next pollutant of interest examined is PM₁₀, as Table III shows its changes for the previous decade. Concerning the concentrations of this air pollutant, there is also a general tendency to reduce its emissions in Norwegian cities. A positive fact is that since 2016 all of the studied cities have managed to maintain an annual mean with the normative limit of 25 $\mu\text{g}/\text{m}^3$. However, in the case of Oslo, it seems that the levels of PM₁₀ are being kept the same for almost the whole period, very close to the indicated limit. In the last years, a particularly impressive improvement in reducing the concentration of this air pollutant is documented in Stavanger.

Concerning PM_{2.5} concentrations, it can be claimed that the implemented measures to keep the levels below annual means of 15 $\mu\text{g}/\text{m}^3$ have proven to be successful, as an exceeding of this limit is not observed for the last decade (Table IV). There is a general decline in emissions, though, in the case of Oslo, this seems harder to state since the value has been fluctuating. This is the only city in which the target of 8 $\mu\text{g}/\text{m}^3$ has not been achieved yet.

C. Air pollution and weather-based variables

In addition to the primary relationship of interest between population densities and air pollution, we also explored the relationship between the latter and the weather-based variables we retrieved from the publicly available sources. This was motivated by the literature review presented in section 2.2 and the ambition of how influential the weather conditions are compared to a socio-demographic factor such as population density. Table VI illustrates the results of the run regression models for available weather data as the highlighted value indicates r^2 and the italic number below represents Pearson's coefficient (r) for each correlation, see Table V.

The presented results indicate some interesting notions with respect to estimated correlations. To simplify them, we can identify whether factors with a homogenous (although not equal) impact on all of the pollutants and weather factors significantly stronger on a particular contaminant. The former group seems to include temperature, atmospheric pressure, and humidity. The latter two weather conditions have a strong negative correlation with all pollutants, meaning the greater

TABLE II
THE ANNUAL MEAN (AVERAGE ANNUAL CONCENTRATIONS) OF NO₂ ($\mu\text{g}/\text{m}^3$) FOR CITIES OF INTEREST FROM 2010 TO 2023

City	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Oslo	62	54	60	61	50	51	56	41	40	39	33	32	29	29
Bergen	55	38	43	41	42	38	41	35	36	32	24	27	24	24
Trondheim	52	46	40	37	35	32	33	28	31	29	22	24	21	19
Stavanger	-	52	45	44	41	37	33	32	28	25	22	27	21	21

TABLE III
THE ANNUAL MEAN (AVERAGE ANNUAL CONCENTRATIONS) OF PM₁₀ ($\mu\text{g}/\text{m}^3$) FOR CASE STUDIES FROM 2010 TO 2023

City	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Oslo	24	22	21	27	25	24	24	22	24	22	23	32	29	29
Bergen	26	-	18	22	19	16	17	14	17	14	15	27	24	24
Trondheim	28	30	29	24	22	16	13	14	17	14	14	24	21	19
Stavanger	29	27	26	28	24	28	22	13	14	13	16	27	21	21

TABLE IV
THE ANNUAL MEAN (AVERAGE ANNUAL CONCENTRATIONS) OF PM_{2.5} ($\mu\text{g}/\text{m}^3$) FOR CASE STUDIES FROM 2010 TO 2023

City	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Oslo	14	14	12	15	11	9	10	9	12	9	12	9	9	7
Bergen	14	9	9	9	9	8	8	7	7	7	8	8	7	7
Trondheim	11	10	10	10	10	7	6	5	7	6	8	7	6	6
Stavanger	12	10	8	10	10	10	10	8	9	9	8	10	7	7

TABLE V
CORRELATIONS BETWEEN POPD AND AIR POLLUTANTS

Correlations	NO ₂ ($\mu\text{g}/\text{m}^3$)	PM _{2.5} ($\mu\text{g}/\text{m}^3$)	PM ₁₀ ($\mu\text{g}/\text{m}^3$)
Pearson's r	0.38	0.42	0.18
r-square (r^2)	0.14	0.18	0.03
Coefficient of covariance (CV)	0.43	0.10	0.01
p-value (standard error)	< 0.01	0.013	0.26

their values are, the lower the concentration of air pollutants will be. However, the average temperature does not seem to influence the concentrations of NO₂, PM_{2.5}, and PM₁₀.

The weather factors, such as precipitation, sunshine hours, and wind, seem to affect the three pollutants differently. For instance, the level of precipitation (both in terms of the average mean of absolute value or as an average number of high precipitation days – above 1.0 mm) influence to a greater extent the concentrations of PM₁₀. A little lower but still in a moderate and negative direction, the precipitation impacts PM_{2.5} and NO₂. On the contrary, the wind speed has a stronger negative correlation with the concentration of NO₂, a minor impact on PM_{2.5} and an insignificant correlation with PM₁₀. Lastly, we can outline the significant impact that sunshine hours have on the concentration of PM₁₀; the data suggest a more moderate effect on PM_{2.5} and a minor one on NO₂.

D. Discussion

After presenting the generated outcomes from the undertaken data processing, we can outline some specific empirical assumptions. First of all, it is essential to address the main research question of whether there is a correlation between population densities and air pollution in the examined case, i.e., the municipalities of the four largest Norwegian cities. The

processed data suggests a positive correlation, but this does not seem highly influential compared with other geographic-meteorological factors. Furthermore, an assumption that we might make is that the positive impact of population density on air quality seems to decrease with time. Based on the examined cases, there are minor increases in population density figures for Bergen, Trondheim, and Stavanger (excluding 2020) and notable growth of the variable in the case of Oslo. However, looking at the concentration of the observed air pollutants, there are strong tendencies of reduction. This trend demonstrates that the policies for reducing air pollution by promoting electric cars (and vehicles) in Norway and fossil-fuel-free modes of transport in cities are somewhat successful.

The second significant outcome from the executed study is that specific weather conditions seem to be much more influential to air pollution than the density of people and activities within the explored context. The annual mean of the average temperature is the only weather-based factor estimated with lower significance to air pollution than population densities in the examined cases. Variables such as atmospheric pressure, humidity, precipitation, and sunshine hours demonstrate strong correlations to at least one of the pollutants of interest. The former two have a strong negative correlation with all of NO₂, PM₁₀ and PM_{2.5}. At the same time, precipitation and sunshine hours predominantly influence the two examined

TABLE VI
RESULTS OF THE RUN REGRESSION MODELS FOR AVAILABLE WEATHER DATA AS THE HIGHLIGHTED VALUE INDICATES R2 AND THE ITALIC NUMBER
BELOW REPRESENTS PEARSON'S FOR EACH CORRELATION

	Temperature	Atmospheric Pressure	Humidity	Wind	Precipitation annually	Precipitation days	Sunshine hours
NO2	0.01 (-0.08)	0.95 (-0.97)	0.81 (-0.9)	0.39 (-0.62)	0.06 (-0.24)	0.39 (-0.63)	0.15 (0.39)
PM2.5	0.02 (0.15)	0.83 (-0.91)	0.78 (-0.88)	0.13 (-0.35)	0.16 (-0.39)	0.52 (-0.72)	0.39 (0.55)
PM10	0 (0.02)	0.48 (-0.69)	0.63 (-0.8)	0.04 (-0.21)	0.63 (-0.8)	0.82 (-0.9)	0.78 (0.89)

variants of PM, respectively, negatively (for precipitation) and positively (for sunshine hours).

V. CONCLUSION AND FUTURE RESEARCH

The applied method promotes a simplistic, straightforward approach to combine different types of data, such as weather monitoring parameters, geo-database, and census data. We have succeeded in comparing the impact of various factors upon the same variable, i.e., concentration of air pollutants, demonstrating to us the higher relevance of certain geographic-meteorological factors on air quality compared to population densities. However, the latter's positive correlation still suggests exploring avenues of further research regarding this finding. As a potential subsequent endeavour in this direction, we see more elaborated regression models to integrate both types of factors in a single equation.

Elaborating further on this line of thought, a potent idea can be to look at additional traffic data, and electrical vehicle adoption in these municipalities as well as changes in means of transportation throughout the years. Socio-economic variables (anthropogenic factors) that impact the concentration of air pollutants. This would be a worthwhile effort to develop a more holistic understanding of the human impact on air quality in the cases of interest. Interesting explorations could incorporate aspects such as car ownership per capita, the percentage of electric cars, and the degree of use of other mobility modes.

Lastly, the technical aspect of the data collection can be improved by incorporating an additional number of monitoring stations to effectively document the air pollution emissions in two of the cases. Potential possibilities to enhance this are the employment of low-cost sensors and mobile monitoring stations. When combined with the official monitoring stations and model calculations, air quality data would eventually offer a high spatial resolution for in-depth research of particular cases, e.g., municipality, city, neighbourhood, etc.

REFERENCES

- [1] Julian Allen, Michael Browne, and Tom Cherrett. Investigating relationships between road freight transport, facility location, logistics management and urban form. *Journal of transport geography*, 24:45–57, 2012.
- [2] Maximilian Auffhammer, Solomon M Hsiang, Wolfram Schlenker, and Adam Sobel. Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, 2013.
- [3] Rainald Borck and Philipp Schrauth. Population density and urban air quality. *Regional Science and Urban Economics*, 86:103596, 2021.
- [4] Tone Bruun, Eva Marie-Louise Denison, Linn Gjersing, Trine Husøy, Ann Kristin Skrindo Knudsen, and Bjørn Heine Strand. Public health report—short version health status in norway 2018. 2018.
- [5] Claire E Campbell, Devyn L Cotter, Katherine L Bottenhorn, Elisabeth Burnor, Hedyeh Ahmadi, W James Gauderman, Carlos Cardenas-Iniguez, Daniel Hackman, Rob McConnell, Kiros Berhane, et al. Air pollution and emotional behavior in adolescents across the us. *medRxiv*, 2023.
- [6] Judith C Chow. Measurement methods to determine compliance with ambient air quality standards for suspended particles. *Journal of the Air & Waste Management Association*, 45(5):320–382, 1995.
- [7] Alan Earnshaw and Norman Neill Greenwood. *Chemistry of the Elements*, volume 60. Butterworth-Heinemann Oxford, 1997.
- [8] Marloes Eeftens, Rob Beelen, Kees De Hoogh, Tom Bellander, Giulia Cesaroni, Marta Cirach, Christophe Declercq, Audrius Dedele, Evi Dons, Audrey De Nazelle, et al. Development of land use regression models for pm2. 5, pm2. 5 absorbance, pm10 and pmcoarse in 20 european study areas; results of the escape project. *Environmental science & technology*, 46(20):11195–11205, 2012.
- [9] US EPA. National ambient air quality standards for particulate matter, proposed rule. *Fed Regist.*, 77:38889–39055, 2012.
- [10] Jane Jacobs. *Jane Jacobs. The Death and Life of Great American Cities*, 21(1):13–25, 1961.
- [11] Lu Liang and Peng Gong. Urban and air pollution: a multi-city study of long-term effects of urban landscape patterns on air quality trends. *Scientific reports*, 10(1):18618, 2020.
- [12] Petros Mouzourides, Prashant Kumar, and Marina K-A Neophytou. Assessment of long-term measurements of particulate matter and gaseous pollutants in south-east mediterranean. *Atmospheric Environment*, 107:148–165, 2015.
- [13] Petter Næss. Urban form and travel behavior: Experience from a nordic context. *Journal of Transport and Land use*, 5(2):21–45, 2012.
- [14] United Nations. New urban agenda. In *Habitat III—The United Nations Conference on Housing and Sustainable Urban Development*, 2017.
- [15] Jonathan Norman, Heather L MacLean, and Christopher A Kennedy. Comparing high and low residential density: life-cycle analysis of energy use and greenhouse gas emissions. *Journal of urban planning and development*, 132(1):10–21, 2006.
- [16] Norwegian Environment Agency. *Limit values and national objectives: Proposed long-term health-based national targets and revised limit values for local air quality*. Report M-129-2014. Norwegian Environment Agency, Oslo, 2014. in Norwegian.
- [17] Norwegian Institute of Public Health - NIPH. *Air quality criteria: effects of air pollution on health*. Norwegian Institute of Public Health, Oslo, 2013.
- [18] Norwegian Institute of Public Health - NIPH. *Public Health Report: Health Status in Norway 2018*. Norwegian Institute of Public Health, Oslo, 2018.
- [19] World Health Organization. *Air quality guidelines: global update 2005: particulate matter, ozone, nitrogen dioxide, and sulfur dioxide*. World Health Organization, 2006.
- [20] C Arden Pope III and Douglas W Dockery. Health effects of fine particulate air pollution: lines that connect. *Journal of the air & waste management association*, 56(6):709–742, 2006.
- [21] Xiaoliang Qin, Lujian Hou, Jian Gao, and Shuchun Si. The evaluation and optimization of calibration methods for low-cost particulate matter sensors: Inter-comparison between fixed and mobile methods. *Science of the total environment*, 715:136791, 2020.
- [22] Raymond Unwin. *The Town Extension Plan*. Manchester University Press, 1912.
- [23] Frank Lloyd Wright. *The disappearing city*. 1932.

Secretary problem revisited: Optimal selection strategy for top candidates using one try in a generalized version of the problem

Lubomír Štěpánek^{1, 2, 3}

¹Department of Statistics and Probability

²Department of Mathematics

Faculty of Informatics and Statistics

Prague University of Economics and Business

W. Churchill's square 4, 130 67 Prague, Czech Republic

lubomir.stepanek@vse.cz

&

³Institute of Biophysics and Informatics

First Faculty of Medicine

Charles University

Salmovská 1, Prague, Czech Republic

lubomir.stepanek@lf1.cuni.cz

Abstract—This paper explores a novel variation of the classical secretary problem, commonly referred to as the marriage or best choice problem. In this adaptation, a decision-maker sequentially dates $n \in \mathbb{N}$ candidates, each uniquely ranked without ties from 1 to n . The decision strategy involves a preliminary non-selection phase of the first $d \in \mathbb{N}$ candidates where, $d < n$, following which the decision-maker commits to the first subsequent candidate who surpasses all previously evaluated candidates in quality. The central focus of this study is the derivation and analysis of $P(d, n, k)$, which denotes the probability that the selected candidate, under the prescribed strategy, ranks among the top $k \in \mathbb{N}$ overall candidates, where $k \leq n$. This investigation employs combinatorial probability theory to formulate $P(d, n, k)$ and explores its behavior across various parameter values of d , n , and k . Particularly, we seek to determine in what fraction of the entire decision process should a decision-maker stop the non-selection phase, i.e., we search for the optimal proportion $\frac{d}{n}$, that maximizes the probability $P(d, n, k)$, with a special focus on scenarios where k is in generally low. While for $k = 1$, the problem is simplified to the classical secretary problem with $\frac{d}{n} \approx \frac{1}{e}$, our findings suggest that the strategy's effectiveness is optimized for portion $\frac{d}{n}$ decreasing below $\frac{1}{e}$ as k increases. Also, intuitively, probability $P(d, n, k)$ increases as k increases, since the number of acceptable top candidates increases. These results not only extend the classical secretary problem but also provide strategic insights into decision-making processes involving ranked choices, sequential evaluation, and applications of searching not necessarily the best candidate, but one of the best candidates.

I. INTRODUCTION

THE secretary problem demonstrates a scenario involving optimal stopping theory, which is studied extensively in the fields of applied probability, statistics, and decision theory. Known under various names such as the marriage problem, the sultan's dowry problem, the fussy suitor problem, the googol

game, and the best choice problem, its solution has garnered attention due to the intriguing nature of its simple yet effective decision strategy often referred to as the 37% rule [1].

In the classical form of the problem, an administrator aims to hire the best secretary out of n (uniquely) rankable applicants. Each applicant is interviewed one at a time in random order, with an immediate decision required at the end of each interview. Once rejected, an applicant cannot be recalled. The challenge lies in making a decision with incomplete information about the quality of unseen applicants, necessitating a strategy that balances the risk of passing up the best candidate against the potential for future superior candidates. The odds algorithm provides the shortest rigorous proof for this problem, establishing that the optimal win probability is always at least $\frac{1}{e}$, with the optimal stopping rule being to reject the first $\sim \frac{n}{e}$ applicants, i.e. roughly 37 % of n , and then stop at the first applicant who is better than all previously interviewed candidates.

Other modifications of the secretary problem explore various strategic nuances. One such variant is the “postdoc” problem [2], where the objective shifts from selecting the best to the second-best candidate (because the “best” will go to Harvard). Theoretical analysis shows that the success probability for this variant with an even number of applicants is exactly $\frac{0.25n^2}{n(n-1)}$, which simplifies to approximately 1/4 as n grows large. This change underscores the subtlety needed in planning and execution when the goal is not to secure the top choice but a candidate just slightly less optimal.

Further expanding the range of strategic considerations, another version allows multiple selections, aiming to secure the top- k candidates using k tries [3]. Here, the challenge

grows with k , as each choice potentially affects subsequent selections. Research indicates that the initial non-selection phase should last approximately $\lfloor \frac{n}{ke^{1/k}} \rfloor$ candidates, maximizing the chance of selecting all top- k candidates, which converges to $\frac{1}{ek}$ in probability as n becomes very large. In a sophisticated variant of the secretary problem, a decision-maker is granted multiple attempts to select the best candidate [4], each governed by a distinct set of r decision thresholds (a_1, a_2, \dots, a_r) , where $a_1 > a_2 > \dots > a_r$. As the number of interviews approaches infinity, the threshold for each decision-maker converges to ne^{-k_i} , where k_i is a defined constant [5]. Bruss and Louchard in [6] explored online selection strategies for choosing the κ best objects from n sequentially observed, rankable objects, with a focus on threshold functions that account for past selections and their asymptotic behavior as $n \rightarrow \infty$.

In this work, we extend the classic framework to develop and analyze $P(d, n, k)$, the probability that after skipping first d candidates, the candidate selected using a stopping rule, i.e., the one that got the highest ranking so far, ranks among the top k candidates out of n candidates in total. Our study examines the effects of varying the parameters d , n , and k to determine optimal selection strategy, particularly we search for proportion $\frac{d}{n}$ that maximizes probability $P(d, n, k)$. Besides different decision thresholds, we specifically explore the efficacy and the impact of increasing k on the strategy's performance. Skipping the first d candidates allows for establishing a robust benchmark of candidate quality without prematurely committing, thus striking a critical balance between not selecting too early (if d is small), which risks missing higher-quality candidates appearing later, and not starting too late (when d is large), which risks missing optimal candidates that have already been evaluated. This opens room for optimizing the value of d or $\frac{d}{n}$ which maximizes probability $P(d, n, k)$ of selecting one of the top k candidates.

II. A MODIFIED SECRETARY PROBLEM OF SEARCHING FOR TOP CANDIDATES

The classical secretary problem, also known as the marriage problem, traditionally focuses on identifying the optimal strategy to select the best candidate from a sequentially reviewed set. This section introduces a modified version of the secretary problem that expands the objective to include not just the best candidate but potentially any of the top few candidates, based on predefined criteria. This modification introduces a more complex and realistic scenario that decision-makers often encounter in various practical applications, from hiring processes to academic admissions.

A. Problem setup and notation

In this modified framework, a decision-maker sequentially dates $n \in \mathbb{N}$ candidates, each uniquely ranked from 1 to n without ties. The candidates are reviewed one at a time in a random sequence, and the decision-maker must decide immediately after each interview whether to select the candidate or continue with the process. Once a candidate is rejected,

they cannot be recalled. The decision-maker initially evaluates d candidates without selecting any of them, where $1 \leq d < n$. This phase is crucial for establishing a benchmark against which all subsequent candidates are compared. It allows the decision-maker to gain a clear understanding of the average candidate quality, setting a standard that must be exceeded to initiate the selection phase. Starting from the $(d + 1)$ -th candidate and continuing with $(d + 2)$ -th, $(d + 3)$ -th candidate, \dots , the decision-maker selects the first candidate who surpasses all previously evaluated candidates in quality. This approach aims to maximize the likelihood of choosing one of the top-ranked candidates by ensuring a thorough comparison to a well-established quality benchmark.

B. Description of the strategy

The specifics of the strategy can be further distilled into several key steps.

- (i) Evaluate the first d candidates without selecting any (non-selection phase).
- (ii) Begin the selection process with the $(d + 1)$ -th candidate.
- (iii) Continue with $(d + 2)$ -th, $(d + 3)$ -th candidate, \dots , and stop at the first candidate who is better than all the evaluated candidates so far.
- (iv) If no such candidate is found by the end, either select the last candidate or leave the position unfilled, depending on specific rules which may be predefined in the problem statement.

This modified approach introduces a dynamic element to the decision-making process, where the decision-maker's strategy adapts based on the outcomes of initial evaluations. Let $P(d, n, k)$ be a probability that the selected candidate, using the prescribed strategy, is among the top k ranked candidates. The objective is to maximize the probability $P(d, n, k)$, which quantifies the success of the strategy in terms of selecting a top-ranked candidate. The subsequent section will delve into the solution methods and analytical techniques used to derive and maximize $P(d, n, k)$, offering insights into the optimal values of $\frac{d}{n}$ and the conditions under which the strategy succeeds.

C. Analytical derivation of $P(d, n, k)$

In this section, we delve into the analytical workings of the formula $P(d, n, k)$, which quantifies the probability that the selected candidate is among the top k out of n candidates, following a strategy where the first d candidates are merely evaluated and not selected, $1 \leq d < n$. The derivation involves considering each candidate i , where i ranges from $d + 1$ to n , and calculating the probability that this i -th candidate is one of the k best.

Assume that the candidate i is being considered for selection, with $i \in \{d + 1, d + 2, \dots, n\}$. The strategy entails skipping the initial d candidates, so the analysis starts from the $(d + 1)$ -th candidate. Several aspects ensure that the i -th candidate, who is selected, is one of top k candidates.

- (i) To select i -th candidate, the decision-maker couldn't meet before any candidate rated higher than any of first

d candidates. Thus, the maximum rating among the first $i-1$ candidates (including the skipped d candidates) must lie within these d candidates, occurring with a probability of $\frac{d}{i-1}$.

(ii) The selected i -th candidate must be among the top k candidates in total.

- The probability that the i -th candidate is the absolute best among all n candidates is $\frac{1}{n}$.
- The probability that the i -th candidate is the second-best involves two conditions: first, the i -th position must actually be the second highest, which occurs with probability $\frac{1}{n}$, and second, the best candidate is not among the first i candidates but among those who follow (otherwise, they should be selected as a candidate with the highest ranking so far), which occurs with probability $\frac{n-i}{n-1}$. Therefore, the combined probability is

$$\frac{1}{n} \cdot \frac{n-i}{n-1}.$$

- For the i -th candidate to be the third-best, the logic extends further: the probability of being third is compounded by the likelihood that exactly two candidates ranked higher are still to come after candidate i , so equal to $\frac{\binom{n-i}{2} \cdot 2!}{\binom{n-1}{2} \cdot 2!}$. This probability is calculated as

$$\frac{1}{n} \cdot \frac{\binom{n-i}{2} \cdot 2!}{\binom{n-1}{2} \cdot 2!} = \frac{1}{n} \cdot \frac{(n-i)(n-i-1)}{(n-1)(n-2)}.$$

- Extending this to the general case for the j -th best, where $1 \leq j \leq k$, the probability that the i -th candidate is the j -th best can be similarly modeled. It is the product of the probability that the i -th position is the j -th highest, $\frac{1}{n}$, and that all higher-ranked $j-1$ candidates appear among those $n-i$ yet to be seen, $\frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n-1}{j-1} \cdot (j-1)!}$, so

$$\begin{aligned} \frac{1}{n} \cdot \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n-1}{j-1} \cdot (j-1)!} &= \frac{\binom{n-i}{j-1} \cdot (j-1)!}{n \cdot \frac{(n-1)!}{(j-1)!} \cdot (n-j)! \cdot (j-1)!} \\ &= \frac{\binom{n-i}{j-1} \cdot (j-1)!}{n \cdot \frac{(n-1)!}{(n-j)!}} \\ &= \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\frac{n!}{(n-j)!}} \\ &= \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\frac{n!}{(n-j)! \cdot j!} \cdot j!} \\ &= \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n}{j} \cdot j!} \end{aligned} \quad (1)$$

The i -th selected candidate can be the first, the second, ..., or the k -th best out of all n candidates, so $j \in \{1, 2, \dots, k\}$. Since these states are obviously disjunctive, the probability that i -th selected candidate

(for fixed $i \in \{d+1, d+2, \dots, n\}$) is among top k is, using formula (1), equals to

$$\sum_{j=1}^k \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n}{j} \cdot j!}. \quad (2)$$

Since conditions (i) and (ii) from previous analysis are independent (condition (i) deals with arrangement of first $i-1$ candidates while condition (ii) handles last $n-i+1$ candidates sequence), we can derive that the probability for the i -th candidate being among the top k and better than the previous maximum observed among the first $i-1$ candidates can be represented (for fixed $i \in \{d+1, d+2, \dots, n\}$) as

$$\frac{d}{i-1} \sum_{j=1}^k \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n}{j} \cdot j!}, \quad (3)$$

and since i -th candidate can be $(d+1)$ -th, $(d+2)$ -th, ..., n -th one (which are disjunctive), we get the complete probability that the i -th candidate is among the top k and with the highest ranking so far out of the first $i-1$ candidates, as follows

$$P(d, n, k) \stackrel{(1,2,3)}{=} \sum_{i=d+1}^n \frac{d}{i-1} \sum_{j=1}^k \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n}{j} \cdot j!}. \quad (4)$$

III. OPTIMAL STRATEGY FOR THE MODIFIED SECRETARY PROBLEM OF SEARCHING FOR TOP CANDIDATES

We assume that a decision-maker knows or pre-estimates the number n of all candidates and also decides how many top candidates k are relevant for their selection. However, the decision-maker would like to know, at which candidate to stop only evaluate rating and start possible selecting to maximize probability $P(d, n, k)$ of selecting one of top k candidates. In other words, the decision-maker would like to know the value of d , or value of $\frac{d}{n}$.

A. Maximizing $P(d, n, k)$ for fixed $n \in \mathbb{N}$ and $k = 1$ with respect to $1 \leq d < n$

Let's focus on simplifying the formula for $P(d, n, k)$ when $k = 1$. This specific case indeed reverts the problem to the classical secretary problem [1], where the goal is to maximize the probability of selecting the best candidate out of n . For $k = 1$, formula (4) simplifies significantly,

$$\begin{aligned} P(d, n, 1) &= \sum_{i=d+1}^n \frac{d}{i-1} \sum_{j=1}^1 \frac{\binom{n-i}{j-1} \cdot (j-1)!}{\binom{n}{j} \cdot j!} = \\ &= \sum_{i=d+1}^n \frac{d}{i-1} \frac{\binom{n-i}{1-1} \cdot (1-1)!}{\binom{n}{1} \cdot 1!} = \\ &= \sum_{i=d+1}^n \frac{d}{i-1} \frac{\binom{n-i}{0} \cdot (0)!}{\binom{n}{1} \cdot 1!} = \\ &= \sum_{i=d+1}^n \frac{d}{i-1} \frac{1}{n} = \frac{d}{n} \sum_{i=d+1}^n \frac{1}{i-1} = \\ &= \frac{d}{n} \sum_{i=d}^{n-1} \frac{1}{i}. \end{aligned} \quad (5)$$

Formula (5) can be simplified using harmonic series, where $H_n = \sum_{i=1}^n \frac{1}{i}$. Therefore, the partial sum from d to $n-1$ as in formula (5) can be expressed as

$$\sum_{i=d}^{n-1} \frac{1}{i} = \sum_{i=1}^{n-1} \frac{1}{i} - \sum_{i=1}^{d-1} \frac{1}{i} = H_{n-1} - H_{d-1},$$

thus, we improve formula (5) as

$$P(d, n, 1) = \frac{d}{n} \sum_{i=d}^{n-1} \frac{1}{i} = \frac{d}{n} (H_{n-1} - H_{d-1}). \quad (6)$$

For large n , the harmonic number H_n can be approximated using the natural logarithm as $H_n \approx \ln(n) + \gamma$, where γ is the Euler-Mascheroni constant. Applying this to our partial sum in formula (6), we get

$$\begin{aligned} P(d, n, 1) &= \frac{d}{n} (H_{n-1} - H_{d-1}) \approx \\ &\approx \frac{d}{n} (\ln(n-1) + \gamma - (\ln(d-1) + \gamma)) \approx \\ &\approx \frac{d}{n} (\ln(n-1) - \ln(d-1)) \approx \\ &\approx \frac{d}{n} \cdot \ln\left(\frac{n-1}{d-1}\right). \end{aligned} \quad (7)$$

Let's take a derivative of $P(d, n, 1)$ from formula (7) with respect to d searching for the value of d maximizing $P(d, n, 1)$,

$$\begin{aligned} \frac{\partial}{\partial d} P(d, n, 1) &\approx \frac{\partial}{\partial d} \left\{ \frac{d}{n} \cdot \ln\left(\frac{n-1}{d-1}\right) \right\} = \\ &= \frac{1}{n} \cdot \ln\left(\frac{n-1}{d-1}\right) + \frac{d}{n} \cdot \frac{1}{d-1} \left(-\frac{n-1}{(d-1)^2}\right) = \\ &= \frac{1}{n} \cdot \ln\left(\frac{n-1}{d-1}\right) - \frac{d}{n(d-1)}. \end{aligned} \quad (8)$$

Putting derivative $\frac{\partial}{\partial d} P(d, n, 1)$ from formula (8) equal zero, we get

$$\begin{aligned} \frac{\partial}{\partial d} P(d, n, 1) &\equiv 0 \\ \frac{1}{n} \cdot \ln\left(\frac{n-1}{d-1}\right) - \frac{d}{n(d-1)} &= 0 \\ \frac{1}{n} \cdot \ln\left(\frac{n-1}{d-1}\right) &= \frac{d}{n(d-1)} \\ \ln\left(\frac{n-1}{d-1}\right) &= \frac{d}{d-1}, \end{aligned}$$

and for large n and d is $\frac{n-1}{d-1} \approx \frac{n}{d}$ and $d \approx d-1$, so

$$\begin{aligned} \ln\left(\frac{n-1}{d-1}\right) &= \frac{d}{d-1} \\ \ln\left(\frac{n}{d}\right) &\approx \frac{d}{d} = 1, \end{aligned} \quad (9)$$

which results in $\frac{n}{d} \approx e$ or $\frac{d}{n} \approx \frac{1}{e}$, where e is Euler constant. The well-known 37% rule comes from the equation $\frac{d}{n} \approx \frac{1}{e}$ in formula (9) since $\frac{1}{e} \approx 0.369$.

B. Maximizing $P(d, n, k)$ for fixed $n \in \mathbb{N}$ and $k > 1$ with respect to $1 \leq d < n$

To simplify and analyze $P(d, n, k)$ in a continuous manner which would enable us to search for $\frac{d}{n}$ maximizing $P(d, n, k)$, we consider a transformation using continuous approximations. Assuming a large n , the sums can be approximated by integrals,

$$P(d, n, k) \approx \int_d^n \frac{d}{x-1} \sum_{j=1}^k \left(\frac{1}{j} \cdot \frac{\binom{n-x}{j-1}}{\binom{n}{j}} \right) dx. \quad (10)$$

Given the transcendental nature of the expressions in formula (10), involving exponential and logarithmic functions in integral form, and the eventual use of the gamma function $\Gamma(x) = (x-1)!$ in place of combinatorial numbers, numerical methods are preferred for solving the optimal parameters, highlighting the complexity and non-linearity of the problem.

C. Maximizing $P(d, n, k)$ for fixed $n \in \mathbb{N}$ and $1 \leq k \leq n$ with respect to $1 \leq d < n$ using numerical approaches

Formula (10) indicates that a numerical searching for $\frac{d}{n}$ that maximizes $P(d, n, k)$ for fixed $n \in \mathbb{N}$ and $1 \leq k \leq n$, where d is any value in $1 \leq d < n$, could be more promising than an analytical solution.

We set $n = 20$ and searched numerically for $\frac{d}{n}$ that maximizes $P(d, n, k)$ for $k \in \{1, 2, \dots, 20\}$. We used both formula (4) and also a function `getMyProbability`(n, d, k, m) based on Monte Carlo simulation, see Algorithm 1, that estimates $P(d, n, k)$ probability for $m = 50$ random samples of candidates for each combination of values (d, n, k) . The outcomes of function `getMyProbability`(n, d, k, m) should confirm analytical correctness of formula (4) for $P(d, n, k)$ probability. In Algorithm 1, we set $d < n$ because if $d = n$, the selection phase cannot occur, as it begins with the $(d+1)$ -th candidate. We denote the probability as $\hat{P}(d, n, k)$ in Algorithm 1 instead of $P(d, n, k)$, as the algorithm uses a finite number (m) of simulated repetitions. If Algorithm 1 is repeated t times, yielding $\hat{P}(d, n, k)_\tau$ in its τ -th iteration, the following relationship holds, $P(d, n, k) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \hat{P}(d, n, k)_\tau$.

In Fig. 1, we see the values of $\frac{d}{n}$ for $n = 20$ and $k \in \{1, 2, \dots, 8\}$. While the boxplots come from simulated samples of candidates so that function `getMyProbability`(n, d, k, m) estimated value of $P(d, n, k)$ across all $m = 50$ repetitions for a given combination of (d, n, k) , the blue line shows analytically computed values of $P(d, n, k)$ using formula (4).

As we expected, the proportion $\frac{d}{n}$ decreases as k increases – it is feasible, if more candidates are acceptable for selection (larger k), the more of them are likely in last $n-d$ candidates in a row, thus, d and $\frac{d}{n}$ could diminish. As k increases, the probability $P(d, n, k)$ of selecting a top k candidate rises, which is intuitive since accepting more candidates increases the chances of selection. Optimal values of d^* and $\frac{d^*}{n}$ that maximizes $P(d, n, k)$ for $n = 20$ and $k \in \{1, 2, \dots, 20\}$, as

Algorithm 1: Estimation of probability $P(d, n, k)$ of selecting one of top k candidates

```

1 Function getMyProbability( $n, d, k, m$ ):
   Input :  $n$  (total # of candidates),  $d$  (# of
           candidates to meet without selecting),  $k$ 
           (# of top ranked candidates),  $m$  (# of
           scenarios to simulate)
   Output: estimate of  $P(d, n, k)$  probability
2  $c \leftarrow 0$ ;
3 for  $i \leftarrow 1$  to  $m$  do
4    $sample \leftarrow$  sample integers from 1 to  $n$ 
   without replacement;
5   if  $d < n$  then
6      $j \leftarrow d + 1$ ;
7     while  $j \leq n$  and
            $sample[j] < \max(sample[1 : d])$  do
8        $j \leftarrow j + 1$ ;
9     end
10    if  $j \leq n$  and  $sample[j]$  is in the top  $k$ 
        positions of  $n$  then
11       $c \leftarrow c + 1$ ;
12    end
13  end
14 end
15  $\hat{P}(d, n, k) \leftarrow \frac{c}{m}$ ;
16 return  $\hat{P}(d, n, k)$ ;
17 return

```

TABLE I
OPTIMAL VALUES OF d^* AND RATIOS $\frac{d^*}{n}$ THAT MAXIMIZES PROBABILITY $P(d^*, n, k)$, AND MAXIMUM VALUE OF THE PROBABILITY FOR SELECTING ONE OF THE TOP k CANDIDATES OUT OF $n = 20$. THE PROPORTION $\frac{d^*}{n}$ DECREASES AS k INCREASES, INDICATING A SHIFT TOWARDS EARLIER SELECTION FOR LESS RESTRICTIVE OUTCOMES.

k	d^*	$\frac{d^*}{n}$	$P(d^*, n, k)$
1	7	0.350	0.384
2	6	0.300	0.538
3	5	0.250	0.627
4	4	0.200	0.687
5	4	0.200	0.730
6	3	0.150	0.766
7	3	0.150	0.794
8	3	0.150	0.813
9	2	0.100	0.836
10	2	0.100	0.856
11	2	0.100	0.870
12	2	0.100	0.881
13	1	0.050	0.889
14	1	0.050	0.907
15	1	0.050	0.922
16	1	0.050	0.934
17	1	0.050	0.942
18	1	0.050	0.947
19	1	0.050	0.950
20	1	0.050	0.950

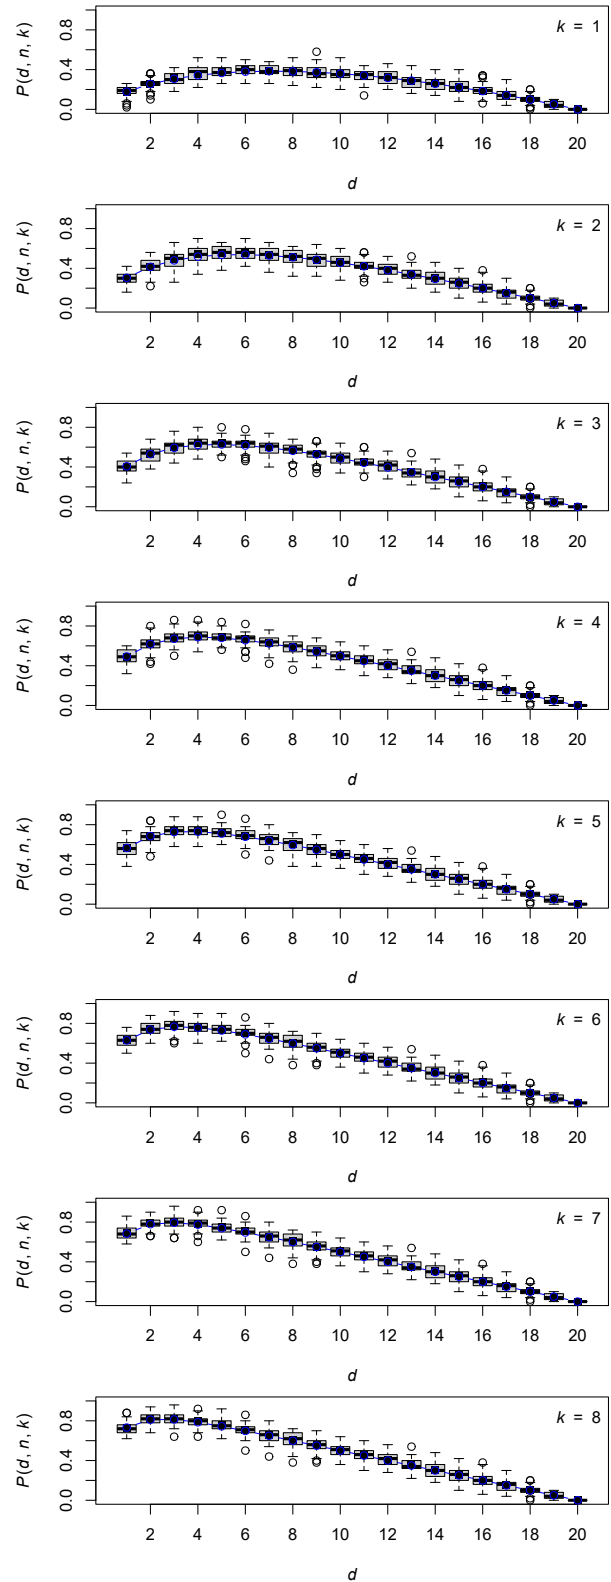


Fig. 1. Optimization of $\frac{d}{n}$ for maximizing $P(d, n, k)$ across k values from 1 to 20 with $n = 20$, using both analytical formula (4) and Monte Carlo simulations via `getMyProbability(n, d, k, m)` (Algorithm 1). As k increases, $P(d, n, k)$ rises while $\frac{d}{n}$ decreases, demonstrating a trade-off in selection strategy efficiency. Boxplots come from $m = 50$ random repetitions of sample generation for given combinations of (d, n, k) and confirm the analytical model's predictions. Blue line stands for values of analytically calculated $P(d, n, k)$.

well as the maximized values of probability $P(d^*, n, k)$, are in Table I, and correspond to outcomes from Fig. 1.

Finally, in Fig. 2, there are values of $\frac{d^*}{n}$ that maximize $P(d, n, k)$ for $n = 100$ and $k \in \{1, 2, \dots, 100\}$, and Fig. 3 shows maximized values of $P(d^*, n, k)$ for varying values $k \in \{1, 2, \dots, 100\}$ and $n = 100$. Similarly as before, maximizing proportions $\frac{d^*}{n}$ decrease and maximized probabilities $P(d^*, n, k)$ increase, as k increases. Both Table I and Fig. 2 show that for a restrictive approach when only top $k = 3$ are acceptable for selection, we need to stop the non-selecting phase not earlier than before skipping first 25 % of candidates, i.e., $\frac{d^*}{n} \approx 0.25$, which could be called as a 25% rule.




Fig. 2. Values of $\frac{d^*}{n}$ that maximize $P(d, n, k)$ for $n = 100$ and varying values $k \in \{1, 2, \dots, 100\}$.




Fig. 3. Maximized values of $P(d^*, n, k)$ for $n = 100$ and varying values $k \in \{1, 2, \dots, 100\}$.

IV. CONCLUSION

This study has extended the classical secretary problem by exploring the probability $P(d, n, k)$ of successfully selecting

one of the top k candidates after initially skipping d candidates. Using probability and combinatorics, we illustrated that analytical methods for searching values of proportions $\frac{d}{n}$ maximizing probability $P(d, n, k)$ are complex, thus, we use Monte Carlo simulations to investigate the optimal strategies for different settings of k .

Our findings indicate a clear strategy shift depending on the rank acceptability, k . As k increases, allowing for a less restrictive choice, the optimal $\frac{d}{n}$ ratio decreases, signifying that an earlier selection becomes preferable. For $k = 1$, which reflects the traditional secretary problem, the optimal skipping strategy is $\frac{d}{n} \approx \frac{1}{e} \approx 0.369$, consistent with the well-known theoretical result of skipping the first 37 % of candidates. Also, probability $P(d, n, k)$ rises as k increases – intuitively, the more candidates are acceptable for selection, the higher is the chance for selecting one of them. When considering only the top $k = 3$ candidates for selection, the non-selecting phase should extend through at least the first 25 % of candidates, effectively establishing a 25% rule.

For real-world applications, such as hiring processes or competitive selection scenarios, these insights can guide more nuanced and more practical strategies that balance risk and reward effectively according to the range of acceptable outcomes. Also, in various applications, we do not want necessarily select the very best candidate, but one of top candidates is enough.

V. ACKNOWLEDGMENT

This paper is supported by the grant IG410023 with no. F4/50/2023, which has been provided by the Internal Grant Agency of the Prague University of Economics and Business.

REFERENCES

- [1] F. T. Bruss, "A unified approach to a class of best choice problems with an unknown number of options," *The Annals of Probability*, vol. 12, no. 3, Aug. 1984.
- [2] R. J. Vanderbei, "The postdoc variant of the secretary problem," *Mathematica Applicanda*, vol. 49, no. 1, Dec. 2021.
- [3] Y. Girdhar and G. Dudek, "Optimal online data sampling or how to hire the best secretaries," in *2009 Canadian Conference on Computer and Robot Vision*. IEEE, May 2009.
- [4] J. P. Gilbert and F. Mosteller, "Recognizing the maximum of a sequence," *Journal of the American Statistical Association*, vol. 61, no. 313, p. 35–73, Mar. 1966.
- [5] T. Matsui and K. Ano, "Lower bounds for bruss' odds problem with multiple stoppings," *Mathematics of Operations Research*, vol. 41, no. 2, p. 700–714, May 2016.
- [6] F. T. Bruss and G. Louchard, "Sequential selection of the κ best out of n rankable objects," *Discrete Mathematics & Theoretical Computer Science*, vol. Vol. 18 no. 3, Oct. 2016.

Non-parametric comparison of survival functions with censored data: A computational analysis of exhaustive and Monte Carlo approaches

Lubomír Štěpánek^{1,2,3}, Filip Habarta¹, Ivana Malá¹, Luboš Marek¹

¹Department of Statistics and Probability

²Department of Mathematics

Faculty of Informatics and Statistics

Prague University of Economics and Business

W. Churchill's square 4, 130 67 Prague, Czech Republic

{lubomir.stepanek, filip.habarta, malai, marek}@vse.cz

&

³Institute of Biophysics and Informatics

First Faculty of Medicine

Charles University

Salmovská 1, Prague, Czech Republic

lubomir.stepanek@lf1.cuni.cz

Abstract—Comparison of two survival functions, which describe the probability of not experiencing an event of interest by a given time point in two different groups, is a typical task in survival analysis. There are several well-established methods for comparing survival functions, such as the log-rank test and its variants. However, these methods often come with rigid statistical assumptions. In this work, we introduce a non-parametric alternative for comparing survival functions that is nearly free of assumptions. Unlike the log-rank test, which requires the estimation of hazard functions derived from (or facilitating the derivation of) survival functions and assumes a minimum number of observations to ensure asymptotic properties, our method models all possible scenarios based on observed data. These scenarios include those in which the compared survival functions differ in the same way or even more significantly, thus allowing us to calculate the p -value directly. Individuals in these groups may experience an event of interest at specific time points or may be censored, i.e., they might experience the event outside the observed time points. Focusing on all scenarios where survival probabilities differ at least as much as observed usually requires computationally intensive calculations. Censoring is treated as a form of noise, increasing the range of scenarios that need to be calculated and evaluated. Therefore, to estimate the p -value, we compare a computationally exhaustive approach that computes all possible scenarios in which groups' survival functions differ as observed or more, with a Monte Carlo simulation of these scenarios, alongside a traditional approach based on the log-rank test. Our proposed method reduces the first type error rate, enhancing its utility in studies where robustness against false positives is critical. We also analyze the asymptotic time complexity of both proposed approaches.

I. INTRODUCTION

SURVIVAL analysis encompasses the study of various time-to-event data, not limited to historical focuses such as death or disease occurrence what could be primarily indicated

by its name. This statistical field characterizes when an event of interest happens or, alternatively, when it does not occur which leads to censoring. The dual nature of the data – combining both the timing of events and their occurrence (or non-occurrence) – distinguishes survival analysis from other statistical methods. Typically, we describe time to an event of interest occurrence for a given individual or a group of individuals using a survival function, which is a function that returns a probability of non-experiencing the event of interest before a given time point.

When comparing two distinct time-to-event survival curves from separate groups, the log-rank test is commonly employed to evaluate the differences, as suggested by Mantel (1966) [1]. The use of the log-rank test is particularly challenged when comparing non-crossing time-event survival curves with uneven censoring between the groups. To enhance the test's efficiency and its robustness against assumptions, various modifications have been proposed. Kong (1997) optimized the log-rank test by adjusting the hazard functions [2], while Song et al. (2007) explored covariate matrix decomposition to establish minimal sample sizes that validate the use of the log-rank test [3]. Additionally, weighted observations have been utilized to correct the test's accuracy, with weights typically greater for earlier events which have more numerous observations, as suggested Peto and Peto (1972) [4], Yang and Prentice (2009) [5], and Li (2018) [6].

Our proposed non-parametric method challenges these limitations by avoiding the estimation of hazard functions and the use of weighting schemes. It explores every conceivable scenario, including those affected by censoring, thus offering a comprehensive approach that traditional methods often can-

not handle due to their computational demands and reliance on strict assumptions. By setting potential event time points for individuals observed as censored, as Štěpánek initially suggested in [7] and [8], we remove data noise and improve the estimation of the survival function for each time point, making our method more adaptable to real-world data.

This approach rigorously evaluates whether the observed differences in survival functions are statistically significant, thereby reducing false positive rates and enhancing methodological robustness. We have developed and applied two computational strategies to estimate p -values: a detailed, computationally intensive approach, and a Monte Carlo simulation, both designed to handle the complexity of evaluating multiple scenarios. For both approaches, we analyzed their asymptotic time complexities.

The manuscript firstly revisits the principles of the log-rank test and its limitations, then introduces our method which involves intensive computational modeling of all potential event scenarios, including censored cases. We provide a relatively detailed analysis of the proposed method's asymptotic time complexity, particularly its p -value calculation using the computationally exhaustive approach and Monte Carlo estimation. The effectiveness of our method is demonstrated through simulation studies comparing p -values calculated using both the log-rank test and our proposed approaches, focusing on their ability to maintain low first type error rates.

II. TRADITIONAL METHODS FOR COMPARISON OF SURVIVAL FUNCTIONS

Firstly, we revisit the log-rank test, examining its principles, assumptions, and limitations.

A. Foundations, assumptions, and limitations of the log-rank test

Foundations of the log-rank test. Assume there are k distinct time points where an event of interest could occur, denoted as t_j for $j \in \{1, 2, 3, \dots, k\}$, and arranged in an ordered tuple $(t_1, t_2, \dots, t_k)^T$. Consider two groups of subjects, labeled as group 1 and 2, $g \in \{1, 2\}$. At each time point t_j , there are $r_{1,j}$ and $r_{2,j}$ individuals at risk (who have not yet experienced the event or have been censored) in groups 1 and 2, respectively, and $d_{1,j}$ and $d_{2,j}$ individuals in each group who have experienced the event. This setup leads to the construction of the contingency table in Table I.

TABLE I
NUMBERS OF INDIVIDUALS EXPERIENCING THE EVENT OF INTEREST IN BOTH GROUPS (1 AND 2) AT TIME POINT t_j .

group	event of interest at time t_j		total
	yes	no	
1	$d_{1,j}$	$r_{1,j} - d_{1,j}$	$r_{1,j}$
2	$d_{2,j}$	$r_{2,j} - d_{2,j}$	$r_{2,j}$
total	d_j	$r_j - d_j$	r_j

The log-rank test evaluates the null hypothesis H_0 that both groups experience identical event rates over time, conditional on fixed rates in the past being the same. Under H_0 , the

observed numbers of events, $D_{1,j}$ and $D_{2,j}$, are modeled as random variables following a hypergeometric distribution with parameters $(r_j, r_{g,j}, d_j)$ for both $g \in \{1, 2\}$. The expected value of $D_{g,j}$ is $\mathbb{E}(D_{g,j}) = r_{g,j} \frac{d_j}{r_j}$, and the variance is $\text{var}(D_{g,j}) = \frac{r_{1,j} r_{2,j} d_j}{r_j^2} \left(\frac{r_j - d_j}{r_j - 1} \right)$. We then compare the observed numbers of events, $d_{g,j}$, for all j , to their expected values. The test statistic for both groups is computed as follows,

$$\chi_{\text{log-rank}}^2 = \frac{\left(\sum_{j=1}^k (d_{g,j} - \mathbb{E}(D_{g,j})) \right)^2}{\sum_{j=1}^k \text{var}(D_{g,j})} \sim \chi^2(1),$$

where $g \in \{1, 2\}$. Under H_0 , the test statistic follows a χ^2 distribution with one degree of freedom. For sufficiently large r_j (at least 30), $\sqrt{\chi_{\text{log-rank}}^2}$ approximates a standard normal distribution. Since $\chi_{\text{log-rank}}^2 \sim \chi^2(1)$, the test statistic can be uniquely transformed into a p -value, representing the conditional probability of observing a test statistic as extreme as or more extreme than the one observed, assuming H_0 is true.

Assumptions and limitations of the log-rank test. It is crucial that right censoring does not differentially affect the event occurrences in both groups. The proportions of censored observations should be nearly equal in both groups; otherwise, the test statistic $\chi_{\text{log-rank}}^2$ could be biased for either group.

Moreover, both the initial total number of individuals at risk and the initial number not experiencing the event should be large enough to meet the Cochran criteria for minimal sample size for χ^2 tests. If these conditions are not met, the $\chi_{\text{log-rank}}^2$ statistic may not fulfill its asymptotic properties, making the estimate numerically unstable.

Additionally, the robustness and statistical power of the log-rank test can be compromised if the proportions $\frac{r_{1,j}}{r_j}$ and $\frac{r_{2,j}}{r_j}$ are not constant across all time points. Significant changes in the survival curves' trends, mutual distances, or crossings can decrease the test's power, making it less likely to reject H_0 when the survival curves actually differ.

III. THE PROPOSED METHOD FOR COMPARISON OF SURVIVAL FUNCTIONS

This section explores our non-parametric approach, which examines all conceivable scenarios where an event could occur at different feasible times for each individual, including those unobserved due to censoring. By adopting all time points for events that come from observed data and calculating a number of all possible scenarios how the events could be registered by individuals in time, we estimate the survival function based on the proportion of individuals who have not experienced the event. We then assess whether the survival functions of the compared groups statistically differ by evaluating the sum of group-based mutual differences across all time points. Theoretically, the proportion of scenarios where the differences are as large or larger than observed corresponds to the p -value, indicating the probability of these findings under the null hypothesis that the survival functions are equivalent. Through this approach, we calculate a range of p -values and

determine whether to reject the null hypothesis based on a predefined confidence level, thus assessing if the survival functions statistically differ.

A. Foundations of the proposed method for comparison of survival functions

We assume two groups of individuals we want to compare so that for each of them we know the same amount of information, coming from data on input. We assume $k \in \mathbb{N}$ distinct time points, denoted t_1, t_2, \dots, t_k . For each individual of both groups, we know two pieces of information (Y_i, \mathcal{E}_i) , where $Y_i \in \{t_1, t_2, \dots, t_k\}$ and $\mathcal{E}_i \in \{\text{event occurred, event did not occur}\}$. While Y_i indicates in which time point an event of interest happened to i -th individual, \mathcal{E}_i describes if it was in fact the event of interest, or rather censoring. Having such information for each individual in a group of $n \in \mathbb{N}$ individuals at all, we can transform pairs (Y_i, \mathcal{E}_i) for $\forall i \in \{1, 2, \dots, n\}$ into grid as in Fig. 1. In lines of the grid in Fig. 1, there are individuals ordered from the one who experienced the event of interest or censoring as first. Thus, if $\Upsilon = (\Upsilon_1, \Upsilon_2, \dots, \Upsilon_n)^T = (\tau_1, \tau_2, \dots, \tau_n)^T$ where $\forall \tau_i \in \{t_1, t_2, \dots, t_k\}$, it is $\tau_1 \leq \tau_2 \leq \dots \leq \tau_n$. While the black dots stand for time points where individuals have not register the event yet including the last point in a row, when the event is experienced, gray dots indicate that an individual was censored at some time point, thus, in theory, the gray dots could be in changed in black if we would know when the event happened (when individual is not censored). White squares stand for time points where an individual experienced neither the event nor censoring.




Fig. 1. The grid of k time points displays n individuals ordered by the timing of their event of interest or censoring. Black dots represent time points where the event has not yet occurred up to the last black point in a row standing for the event, while gray dots indicate censoring, suggesting that these could become black if the event timing were known. White squares denote time points where neither the event nor censoring has occurred.

We want to test the following null hypothesis H_0 against the alternative H_1 as follows,

- H_0 : survival functions do not differ between the groups,
- H_1 : survival functions differ between the groups.

For purposes of statistical inference, we need to calculate p -value as a probability of observing data the same way or

even more contrary the hypothesis H_0 that the survival functions do not differ between the groups. Addressing the p -value calculation in a non-parametric fashion, we have to calculate a number of all scenarios that favor the p -value's definition meaning. The number of scenarios in contradiction to H_0 depends on when censoring happen to censored individuals.

Let $\mathcal{C} \subseteq \{1, 2, \dots, n\}$ be a subset of individuals' indices that have been censored. Considering the censoring arrangement, for each such arrangement of values $\tau_i \in \{t_1, t_2, \dots, t_k\}$ where $i \in \mathcal{C}$, we can calculate a unique p -value, since the entry grids as in Fig. 1 differ as τ_i vary. To do this, we need updated times of events with respect to censoring for all individuals, denoted as $\Upsilon' = (\Upsilon'_1, \Upsilon'_2, \dots, \Upsilon'_n)^T = (\tau'_1, \tau'_2, \dots, \tau'_n)^T$ where

$$\begin{cases} \tau'_i \geq \tau_i, & \forall i \in \mathcal{C}, \\ \tau'_i = \tau_i, & \forall i \in \{1, 2, \dots, n\} \setminus \mathcal{C}. \end{cases} \quad (1)$$

As an illustration, we can compare Fig. 2 and Fig. 3. In both figures, there is $n = 6$ and time points $t \in \{1, 2, \dots, 8\}$. Using previous notations, obviously it is $(\tau_1, \tau_2, \tau_3, \tau_4, \tau_5, \tau_6)^T = (1, 1, 2, 2, 3, 4)^T$ and $\mathcal{C} = (2, 4)^T$ in Fig. 2. In Fig. 3, we set time points of possible event registrations for censored individuals as follows, $\tau_2 = 2$ and $\tau_4 = 6$, thus, $(\tau'_1, \tau'_2, \tau'_3, \tau'_4, \tau'_5, \tau'_6)^T = (1, 2, 2, 6, 3, 4)^T$.




Fig. 2. A grid for an initial dataset with $n = 6$ individuals across $t \in \{1, 2, \dots, 8\}$ time points, where $(\tau_1, \tau_2, \tau_3, \tau_4, \tau_5, \tau_6)^T = (1, 1, 2, 2, 3, 4)^T$ and censored times $\mathcal{C} = (2, 4)^T$.




Fig. 3. A grid for the initial dataset with $n = 6$ individuals across $t \in \{1, 2, \dots, 8\}$ time points from Fig. 2, where the time points for censored events are adjusted, setting $\tau_2 = 2$ and $\tau_4 = 6$ and getting $(\tau'_1, \tau'_2, \tau'_3, \tau'_4, \tau'_5, \tau'_6)^T = (1, 2, 2, 6, 3, 4)^T$, to hypothesize the potential event registrations if not censored.

Once the censoring is arranged and we set Υ'_i for $\forall i \in \{1, 2, \dots, n\}$, we can enumerate a survival function for the

given group of n individuals (and the given censoring arrangement) using an expected survival $\mathbb{E}(S')$ for the group, i.e., a "surface" below a curve enveloping black dots in grid as we defined before and plot in Fig. 1. To be more specific,

$$\hat{\mathbb{E}}(S') = \frac{1}{n} \sum_{i=1}^n \tau'_i, \quad (2)$$

using the same notation as so far. Let's assume two groups with indices 1 and 2 and observed expected survivals, following formula (2), as $\hat{\mathbb{E}}(S'_1)$ and $\hat{\mathbb{E}}(S'_2)$. Now, we can finally calculate p -value, conditional of the given censoring and enabling any further statistical inference, as follows,

$$\begin{aligned} p\text{-value}_{\text{censoring}} &= \\ &= \{p\text{-value} \mid \text{censoring}\} = \\ &= P \left(\begin{array}{l} \text{getting data at least as extreme} \\ \text{as the observed} \end{array} \middle| H_0, \text{censoring} \right) = \\ &= P \left(|\mathbb{E}(S_1) - \mathbb{E}(S_2)| \geq |\hat{\mathbb{E}}(S_1) - \hat{\mathbb{E}}(S_2)| \middle| \text{censoring} \right) = \\ &= P \left(|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)| \right), \end{aligned} \quad (3)$$

thus, in other words, we calculate p -value for a given censoring, denoted as $p\text{-value}_{\text{censoring}}$, as a probability of observing a scenario with an absolute difference between expected survivals as least as extreme as the observed absolute difference of survivals. Since censoring arrangements may differ, we get various p -values from formula (3) stored in vector $\mathbf{p}\text{-value}_{\text{censoring}}$, that may create an interval of p -values as follows,

$$p\text{-value} \in \left\langle \min\{\mathbf{p}\text{-value}_{\text{censoring}}\}, \max\{\mathbf{p}\text{-value}_{\text{censoring}}\} \right\rangle. \quad (4)$$

Assuming a confidence level $\alpha \in (0, 1)$ as an acceptable probability of first type error rate, i.e., false rejection of null hypothesis H_0 which is in fact true, there are two cases crucial for statistical inference, either

$$\max\{\mathbf{p}\text{-value}_{\text{censoring}}\} \leq \alpha,$$

resulting in a strong null hypothesis H_0 rejection, and

$$\min\{\mathbf{p}\text{-value}_{\text{censoring}}\} \leq \alpha,$$

leading to a weak null hypothesis H_0 rejection.

B. Approaches to p -value calculation for the proposed method of survival functions' comparison

In this section, we introduce a computationally exhaustive approach and Monte Carlo simulation-based approach on calculation of p -value following formula (3), and discuss their asymptotic time complexity.

Computationally exhaustive approach for p -value calculation. Within the computationally intensive approach, we work out formula (3) for p -value calculation. Firstly, we realize that the censoring arrangement for an entry grid of

n individuals and k time points is fully defined by $\mathbf{Y}' = (\mathbf{Y}'_1, \mathbf{Y}'_2, \dots, \mathbf{Y}'_n)^T = (\tau'_1, \tau'_2, \dots, \tau'_n)^T$ as comes from formula (1). Assuming the total numbers of all scenarios for both groups, including all possible censoring arrangements \mathbf{Y}'_1 and \mathbf{Y}'_2 , are on grids $n_1 \times k$ and $n_2 \times k$ finite, respectively, we can work out formula (3) as follows,

$$\begin{aligned} p\text{-value}_{\text{censoring}} &= \\ &= \{p\text{-value} \mid \text{censoring}\} = \\ &= P \left(\begin{array}{l} \text{getting data at least as extreme} \\ \text{as the observed} \end{array} \middle| H_0, \text{censoring} \right) = \\ &= P \left(\begin{array}{l} \text{getting data at least as extreme} \\ \text{as the observed} \end{array} \middle| H_0, \mathbf{Y}'_1, \mathbf{Y}'_2 \right) = \\ &= P \left(|\mathbb{E}(S_1) - \mathbb{E}(S_2)| \geq |\hat{\mathbb{E}}(S_1) - \hat{\mathbb{E}}(S_2)| \middle| \mathbf{Y}'_1, \mathbf{Y}'_2 \right) = \\ &= P \left(|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)| \right) = \\ &= \frac{1}{|\{\mathbf{S}_1\}| \cdot |\{\mathbf{S}_2\}|} \cdot \\ &\quad \sum_{\forall s \in \{\mathbf{S}_1\}} \sum_{\forall \sigma \in \{\mathbf{S}_2\}} \mathbb{1}_{\{|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|\}}, \end{aligned} \quad (5)$$

where $\{\mathbf{S}_1\}$ and $\{\mathbf{S}_2\}$ are sets of all possible scenarios for first and second group, respectively, and $\mathbb{1}_{\{\beta\}}$ is an identifier function, so

$$\mathbb{1}_{\{\beta\}} = \begin{cases} 1, & \text{if } \beta \text{ is true,} \\ 0, & \text{if } \beta \text{ is false.} \end{cases} \quad (6)$$

The numbers of all scenarios, $|\{\mathbf{S}_1\}|$ and $|\{\mathbf{S}_2\}|$ are straightforward and can be assessed using stars-and-bars theorem. Assuming grids containing n_1 and n_2 individuals, respectively, and k time points as in Fig. 1, and a fact that for increasing index i , time $\tau_i \in \{0, 1, \dots, k\}$ of event or censoring forms a non-decreasing sequence as in Fig. 3, it is

$$|\{\mathbf{S}_g\}| = \binom{n_g + k}{k} = \binom{n_g + k}{n_g} = \frac{(n_g + k)!}{n_g! k!}, \quad (7)$$

where $g \in \{1, 2\}$. Formula (7) enables us to investigate asymptotic time complexity of p -value from formula (5). As we can see in Algorithm 1, total counts of scenarios $|\{\mathbf{S}_1\}|$ and $|\{\mathbf{S}_2\}|$ (in the fraction part of formula (5)) are calculated asynchronously (line 1 in Algorithm 1), so an asymptotic time complexity of their calculation with respect to formula (7), using Bachmann–Landau logic [9] and unit time steps for basic arithmetic operations, is

$$\Theta(n_g + k - 1 + n_g - 1 + k - 1) \approx \Theta(2n_g + 2k) \approx \Theta(n_g + k), \quad (8)$$

thus, in total, the fraction part of formula (5) has got asymptotic time complexity $\Theta(\dagger)$ where

$$\Theta(\dagger) \stackrel{(8)}{\approx} \Theta(n_1 + k) + \Theta(n_2 + k) \approx \Theta(n_1 + n_2 + 2k). \quad (9)$$

On the other hand, the count of scenarios with absolute difference of expected survivals greater than or equal to the

observed difference is examined in the summation part of formula (5) exhaustively step by step, considering the total number of combinations of scenarios equal to $|\{\mathcal{S}_1\}| \cdot |\{\mathcal{S}_2\}|$ (lines 2–9 in Algorithm 1). Within each step, condition $|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|$ is checked (lines 5–7 in Algorithm 1) – while the part $|\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|$ is once pre-calculated (so we can ignore its complexity within the loop of steps), the difference $|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)|$ takes $n_1 + n_2$ unit times per each check, i.e., per each step, as comes from formula (2). Thus, asymptotic time complexity $\Theta(\ddagger)$ of the summation part of formula (5) is

$$\begin{aligned}
\Theta(\ddagger) &\stackrel{(7)}{\approx} \Theta \left\{ \binom{n_1+k}{k} \binom{n_2+k}{k} (n_1+n_2) \right\} \leq \\
&\leq \Theta \left\{ \binom{n_1+k}{\frac{n_1+k}{2}} \binom{n_2+k}{\frac{n_2+k}{2}} (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ \frac{(n_1+k)!}{\frac{n_1+k}{2}! \frac{n_1+k}{2}!} \frac{(n_2+k)!}{\frac{n_2+k}{2}! \frac{n_2+k}{2}!} (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ \frac{(n_1+k)! (n_2+k)!}{\left(\frac{n_1+k}{2}!\right)^2 \left(\frac{n_2+k}{2}!\right)^2} (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ \frac{\left(\frac{n_1+k}{2}\right)^{n_1+k} \left(\frac{n_2+k}{2}\right)^{n_2+k}}{\left(\left(\frac{n_1+k}{4}\right)^{\frac{n_1+k}{2}}\right)^2 \left(\left(\frac{n_2+k}{4}\right)^{\frac{n_2+k}{2}}\right)^2} (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ \frac{\left(\frac{n_1+k}{2}\right)^{n_1+k} \left(\frac{n_2+k}{2}\right)^{n_2+k}}{\left(\frac{n_1+k}{4}\right)^{n_1+k} \left(\frac{n_2+k}{4}\right)^{n_2+k}} (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ 2^{n_1+k} \cdot 2^{n_2+k} \cdot (n_1+n_2) \right\} \approx \\
&\approx \Theta \left\{ 2^{n_1+n_2+2k} \cdot (n_1+n_2) \right\}. \tag{10}
\end{aligned}$$

Algorithm 1: Calculation of $p\text{-value}_{\text{censoring}}$ using formula (5)

Result: Calculation of $p\text{-value}_{\text{censoring}}$ based on scenario sets \mathcal{S}_1 and \mathcal{S}_2

```

1 calculate  $|\{\mathcal{S}_1\}|$  and  $|\{\mathcal{S}_2\}|$  using formula (7);
2 initialize count  $c = 0$ ;
3 forall  $s \in \{\mathcal{S}_1\}$  do
4   forall  $\sigma \in \{\mathcal{S}_2\}$  do
5     if  $|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|$  then
6        $c = c + 1$ ;
7     end
8   end
9 end
10  $p\text{-value}_{\text{censoring}} = \frac{c}{|\{\mathcal{S}_1\}| \cdot |\{\mathcal{S}_2\}|}$ ;
11 return  $p\text{-value}_{\text{censoring}}$ 

```

Putting things together, calculation $p\text{-value}_{\text{censoring}}$ from formula (5) takes the following time $\Theta(\bullet)$,

$$\begin{aligned}
\Theta(\bullet) &\approx \Theta(\dagger) + \Theta(\ddagger) \stackrel{(9,10)}{\approx} \\
&\stackrel{(9,10)}{\approx} \Theta \left\{ n_1 + n_2 + 2k + \binom{n_1+k}{k} \binom{n_2+k}{k} (n_1+n_2) \right\}, \tag{11}
\end{aligned}$$

and also

$$\begin{aligned}
\Theta(\bullet) &\approx \Theta(\dagger) + \Theta(\ddagger) \stackrel{(9,10)}{\leq} \\
&\stackrel{(9,10)}{\leq} \Theta \left\{ n_1 + n_2 + 2k + 2^{n_1+n_2+2k} \cdot (n_1+n_2) \right\}. \tag{12}
\end{aligned}$$

Considering formula (1) a $|\mathcal{C}|$ individuals that are censored, in theory, in any time point $\{t_1, t_2, \dots, t_k\}$, the calculation of $p\text{-value}_{\text{censoring}}$ using formula (5) might be repeated at maximum $k^{|\mathcal{C}|}$, thus the asymptotic time complexity could be at maximally

$$\Theta(\bullet) \stackrel{(12)}{\leq} k^{|\mathcal{C}|} \cdot \Theta \left\{ n_1 + n_2 + 2k + 2^{n_1+n_2+2k} \cdot (n_1+n_2) \right\}. \tag{13}$$

Monte Carlo approach for p -value calculation. Using formula (5), we cannot necessary consider every possible scenario, but can randomly select a subset \mathcal{M} of a joint scenarios' set $\{\mathcal{S}_1 \cup \mathcal{S}_2\}$, so $\mathcal{M} \subseteq \{\mathcal{S}_1 \cup \mathcal{S}_2\}$. Then, assuming the same conditions, p -value considering a given censoring setting can be estimated similarly as in formula (5),

$$\hat{p}\text{-value}_{\text{censoring}} = \frac{1}{|\mathcal{M}|} \cdot \sum_{\forall m \in \mathcal{M}} \mathbb{1}_{\{|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|\}}, \tag{14}$$

using the same mathematical notation as above. Investigating Algorithm 2, asymptotic time complexity of Monte Carlo approach could be straightforwardly estimated. The condition $|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|$, which takes $n_1 + n_2$ time units, is checked \mathcal{M} times, thus,

$$\Theta(\bullet) \approx |\mathcal{M}| \cdot (n_1 + n_2), \tag{15}$$

and considering the censoring, it is

$$\Theta(\bullet) \leq k^{|\mathcal{C}|} \cdot |\mathcal{M}| \cdot (n_1 + n_2). \tag{16}$$

IV. SIMULATION STUDY

In this simulation study, we compared the traditional log-rank test with our proposed method, using both exhaustive and Monte Carlo approaches for p -value calculation. The methods were tested on numerous pairs of survival curves assumed to be equivalent to assess the first type error rate, i.e., the frequency of falsely rejecting the null hypothesis. Since the proposed method is non-parametric and robust, we focused on evaluating its first type error rate rather than its power. All computations were performed in R statistical language [10].

Algorithm 2: Calculation of \hat{p} -value_{censoring} using formula (14)

Result: Estimate of \hat{p} -value_{censoring} using a Monte Carlo approach based on a subset \mathcal{M} of scenarios

```

1  $\mathcal{M} \leftarrow$  randomly select a subset from  $\{S_1 \cup S_2\}$ ;
2 initialize count  $c = 0$ ;
3 forall  $m \in \mathcal{M}$  do
4   if  $|\mathbb{E}(S'_1) - \mathbb{E}(S'_2)| \geq |\hat{\mathbb{E}}(S'_1) - \hat{\mathbb{E}}(S'_2)|$  then
5      $c = c + 1$ ;
6   end
7 end
8  $\hat{p}$ -valuecensoring =  $\frac{c}{|\mathcal{M}|}$ ;
9 return  $\hat{p}$ -valuecensoring

```

We generated pairs of groups of varying size of $n \in \{10, 11, \dots, 100\}$ individuals so that their survival functions follow a negatively exponential survival function,

$$S(t) = P(T \leq t) = e^{-\lambda t}, \quad (17)$$

where λ varied in range of $(0.04, 0.06)$. Sizes of the groups randomly differed between the repetitions of the simulation. A total of $m = 1000$ pairs of groups with survival functions were generated, and each pair was analyzed under different levels of censoring, set at 10 %, 20 %, 30 %, and 40 %. The occurrence of p -values as intervals, which comes from the present censoring and formula (4), either containing or does not containing the significance level $\alpha = 0.05$ was recorded, summarizing the first type error rates as frequencies when the p -value interval's maximum is below $\alpha = 0.05$, see Table II.

TABLE II
FIRST TYPE ERROR RATES FOR THE LOG-RANK TEST AND PROPOSED METHODS AT VARYING LEVELS OF CENSORING.

censoring level	log-rank test	proposed method	
		exhaustive approach	Monte Carlo
10 %	0.055	0.041	0.040
20 %	0.053	0.039	0.042
30 %	0.052	0.043	0.041
40 %	0.050	0.038	0.039

Simulations show that the first type error rate is consistently lower for the proposed method, whether using exhaustive or Monte Carlo p -value calculations. This reduction persists

across all levels of censoring, highlighting the method's robustness compared to the log-rank test, especially at higher censoring levels.

V. CONCLUSION

This study introduces a novel, assumption-minimal, non-parametric method for comparing survival functions. Utilizing computationally exhaustive and Monte Carlo simulations for p -value calculation, the method consistently shows lower first type error rates than the log-rank test across various levels of censoring (10 % to 40 %). While the approach involves high asymptotic time complexity during p -value estimation, especially with exhaustive calculations, its reduced first type error rate offers an alternative for survival data analysis, potentially suitable for integration into statistical software.

VI. ACKNOWLEDGEMENT

This research is supported by grant F4/50/2023 from the Internal Grant Agency of the Prague University of Economics and Business.

REFERENCES

- [1] N. Mantel, "Evaluation of survival data and two new rank order statistics arising in its consideration," *Cancer Chemotherapy Reports*, vol. 50, no. 3, pp. 163–170, 1966.
- [2] F. Kong, "Robust covariate-adjusted logrank tests," *Biometrika*, vol. 84, no. 4, pp. 847–862, Dec. 1997.
- [3] R. Song, M. R. Kosorok, and J. Cai, "Robust covariate-adjusted log-rank statistics and corresponding sample size formula for recurrent events data," *Biometrics*, vol. 64, no. 3, pp. 741–750, Dec. 2007.
- [4] R. Peto and J. Peto, "Asymptotically efficient rank invariant test procedures," *Journal of the Royal Statistical Society. Series A (General)*, vol. 135, no. 2, p. 185, 1972.
- [5] S. Yang and R. Prentice, "Improved logrank-type tests for survival data using adaptive weights," *Biometrics*, vol. 66, no. 1, pp. 30–38, Apr. 2009.
- [6] C. Li, "Doubly robust weighted log-rank tests and renyi-type tests under non-random treatment assignment and dependent censoring," *Statistical Methods in Medical Research*, vol. 28, no. 9, pp. 2649–2664, Jul. 2018.
- [7] L. Štěpánek, F. Habarta, I. Malá, and L. Marek, "Analysis of asymptotic time complexity of an assumption-free alternative to the log-rank test," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems*, ser. FedCSIS 2020. IEEE, Sep. 2020, p. 453–460.
- [8] —, *Reducing the First-Type Error Rate of the Log-Rank Test: Asymptotic Time Complexity Analysis of An Optimized Test's Alternative*. Springer International Publishing, Dec. 2021, p. 281–302.
- [9] D. E. Knuth, "Big omicron and big omega and big theta," *ACM Sigact News*, vol. 8, no. 2, pp. 18–24, 1976.
- [10] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2023. [Online]. Available: <https://www.R-project.org/>

Predicting Stock Trends Using Common Financial Indicators: A Summary of FedCSIS 2024 Data Science Challenge Held on KnowledgePit.ai Platform

Aleksandar M. Rakićević*¹, Pavle D. Milošević*¹, Ivana T. Dragović*¹, Ana M. Poledica*¹,
Milica M. Zukanović*¹, Andrzej Janusz^{†‡}, Dominik Ślęzak^{†§}

* University of Belgrade - Faculty of Organizational Sciences, Jove Ilića 154, 11000 Belgrade, Serbia
Email: {aleksandar.rakicevic, pavle.milosevic, ivana.dragovic, ana.poledica, milica.zukanovic}@fon.bg.ac.rs

[†]Institute of Informatics, University of Warsaw, Banacha 2, 02-097 Warsaw, Poland
Email: {a.janusz, slszak}@mimuw.edu.pl

[‡]School of Information Systems, Queensland University of Technology, 2 George Street, 4000 Brisbane, Australia
Email: andrzej.janusz@qut.edu.au

[§]QED Software, Mazowiecka 11/49, 00-052 Warsaw, Poland
Email: dominik.slezak@qedsoftware.com

Abstract—Predictive analytics aims to empower finance professionals to make data-driven decisions, anticipate customer behavior, and navigate the complexities of the financial landscape. One of the tasks in this domain is the prediction of stock trend movements. The goal of the FedCSIS 2024 Data Science Challenge was to build such predictive models based on the financial fundamental data. Such models could have a vital role in algorithmic or manual trading, providing trading signals for making decisions about the time and direction of stock trades. We describe the prepared dataset and challenge task. We also summarize the challenge outcomes and provide insights about the most successful machine learning techniques applied.

Index Terms—data science competitions; KnowledgePit.ai platform; stock market data; automatic trading

I. INTRODUCTION

THE STOCK market is a vital indicator of economic health, reflecting the dynamic interaction of investor sentiment, corporate performance, and macroeconomic trends. The global equity market has almost doubled its value for the last decade. According to *statista*¹, the total market capitalization of companies listed on stock exchanges worldwide increased from 65.04 trillion US dollars in 2013 to 111 trillion in 2023. As of December 2023, America's region has the largest equity market share, with NYSE and NASDAQ as the largest stock exchange operators, followed by Asia-Pacific and EMEA.

Stock market prediction typically implies forecasting price, trend, and direction of movement of stocks and stock market indexes. It is considered a rather challenging task, being volatile, stochastic, nonlinear, and influenced by a large

number of factors. Traditionally, the stock market has been analyzed using technical and fundamental analysis. Further, as a sequence of historical data points (e.g., daily, monthly, quarterly) it is frequently modeled using traditional times series approaches such as statistical ARIMA, exponential smoothing (ES), Facebook's Prophet well-known as the industry standard, and its nonlinear extension NeuralProphet [1]. Since 2015, there has been an exponential growth of research papers investigating ML algorithms for stock market predictions [2]. A comprehensive literature review can be found in [3].

Understanding the dynamics of the stock market and forecasting stock markets have been receiving continuous research attention throughout the previous decades [3]. In most research papers, the US stock market indices the Standard & Poor's 500 (S&P 500), NYSE, NASDAQ and DJIA prevailed as the most common data sources [2]. Among the numerous indices that measure stock market performance, S&P 500 holds a prominent position. S&P 500 comprises 500 of the largest publicly traded companies in the United States and is regarded as a strong indicator of the US economy, but it is also a benchmark for global equity markets. Despite geopolitical tensions and an anticipated recession, the S&P 500 has shown a significant increase in value in the last 5 years. In 2023, the index recorded the highest value, closing at 4,769.83. The S&P 500 provides insights into market trends and risk management strategies. It also helps investors construct portfolios based on S&P 500 constituents. The stock movements impact not only individual investors but also influence institutional strategies, government policies, and international investment flows.

FedCSIS 2024 Data Science Challenge aimed at stock trend forecasting of S&P 500 companies. The scope is reduced to 300 companies that have been part of S&P 500 for the

¹www.statista.com/statistics/274490/global-value-of-share-holdings-since-2000/

last 10 years. Due to the availability of financial statements and the long history of data, these publicly traded companies have been particularly appealing for machine learning (ML) research and performance analysis [4]. The prevailing markets and the most famous stock market indexes are surveyed in [5]. The task of stock trend prediction is valuable as it guides investment decisions and trading, risk management, and portfolio optimization. It provides useful insights into economic and market conditions. Even though there is a vast amount of publications about ML in stock market forecasting, this topic remains attractive for scientists and financial professionals.

The paper is organized as follows: Section II reviews the literature on ML algorithms applied in stock market forecasting. Section III summarizes the history of data science challenges held at KnowledgePit.ai. Section IV outlines the challenge's objective, gives details of the prediction problem that was solved, and describes our baseline solution. Section V discusses some insights from the post-competition analysis of submitted solutions. Section VI concludes the paper.

II. RELATED LITERATURE

In numerous ML studies the research focuses on several directions: classical ML, Ensemble Learning (EL), and Deep Learning (DL) [6], [7], [8]. In the last decade, Support Vector Machines (SVM) and Multi-layer Perceptrons (MLP) have been predominant ML algorithms (approximately 30 %) in stock market forecasting, followed by a group of regression algorithms (linear, logistic, and decision trees), Naïve Bayes (NB), k-nearest Neighbors (kNN), etc. [8].

EL techniques, especially Random Forest (RF), Light Gradient Boosting (LGB), and Extreme Gradient Boosting (XGBoost) have also shown promising forecasting outcomes [6]. A comprehensive evaluation of EL for stock market prediction is given in [9]. Due to their high performance in various data science challenges [10], where they consistently outperformed other algorithms, EL techniques also prevailed throughout our FedCSIS 2024 Data Science Challenge, including particularly the submitted solutions that exceeded the baseline.

In recent years, DL has deserved special attention in stock market prediction. This is because of the availability of financial data with a long history, and the fact that stock forecast is influenced by sentiment described by text [6]. The most common DL algorithms for stock market prediction include deep feed-forward, convolutional, and recurrent neural networks (DNN, CNN, RNN), as well as long short-term memory (LSTM), Gated Recurrent Unit (GRU), and bi-directional LSTM [5], [11]. Among all ML/DL algorithms, based on the recent analysis [12], LSTM is the most preferred model for predicting stock price movements. It is followed by classical ML models (SVM, MLP), next to a bigger family of DL models.

Apart from ML algorithms, it is critical to identify features that affect ML performance. According to a research study on feature selection and extraction for stock market prediction from 2011–2022 [13], the techniques most widely used for this




Fig. 1: The share of different industries in the dataset.

purpose in stock market applications are based on correlations, RF, Principal Component Analysis (PCA), and autoencoders.

III. HISTORY OF KNOWLEDGEPI.T.AI CHALLENGES

FedCSIS data science challenges have been held on KnowledgePit.ai platform since 2014 [14], [15]. The topics included recognizing firefighters' activities based on sensor readings (2015) [16], [17], predicting seismic activity in coal mines (2016) [18], [19], video game data analytics (2017–2019) [20], [21], [22], predicting network device workloads (2020) [23], [24], predicting costs of freight forwarding contracts (2022) [25], [26], and detecting cyber-attacks on IoT devices (2023) [27], [28]. These challenges were highly successful, with more than 1,600 participating teams and thousands of solutions reflecting state-of-the-art methods in the fields such as feature extraction [19], [29], time series forecasting [24], [30], and EL-based prediction models [21], [31].

KnowledgePit.ai has evolved along with FedCSIS. Over the years, the platform's goals shifted from smaller projects to becoming a host for international data science challenges. The functionalities offered by the platform have also expanded to facilitate post-competition data analysis [10]. The most prestigious events in recent years were those hosted for industry clients such as Security on Demand (currently DeepSeas) [32], Information Builders [33], EMCA Software [23] or, as in the case of this particular challenge, Yettel.Bank.

IV. TASK AND RESULTS

The challenge focused on predicting trends in the US stock market based on fundamental financial indicators. A unique, hand-crafted dataset has been prepared for this challenge, encompassing quarterly data from financial statement announcements of 300 companies that constitute the S&P 500 index. Fig. 1 illustrates the sectoral distribution of the selected companies. The dataset covers 10 years, from 2014 to 2023.




Fig. 2: Types of financial indicators used in the dataset.




Fig. 3: The class distribution.

Companies were selected based on data availability; those without sufficient data for most of the ten years and those with a high percentage of missing values were omitted. The final dataset contains 10,000 instances. It has been made available to the community to facilitate post-competition research.

Each data instance contains information on the company's sector, values for 58 financial indicators, their 1-year (absolute) changes, target class information (the column 'Class'), and risk-return performance for a period after the announcement (the column 'Perform'). The indicators were chosen based on data availability, literature review, and advice from domain experts. Fig. 2 illustrates the number of indicators within each indicator category. All indicator values are annualized using the Trailing Twelve Months method to neutralize seasonal variation. To prevent participants from gaining an unfair advantage by "looking-ahead", the dataset was anonymized by removing the names of companies and timestamps.

A. Data preparation

The dataset contains two distinct types of missing values with different semantics. NA values indicate that a certain financial indicator does not apply to a company. Empty cells represent conventional missing values (nulls) due to unavailable or missing information. The dataset contains approximately 2.75% NA values and 0.29% null values. The maximum NA percentage for a single attribute is 19.16%. The maximum null value percentage for a single attribute is 5.64%. This presents serious challenges when working with this dataset.

B. Evaluation procedure

In this challenge, participants were asked to solve a three-class classification problem, where:

- **Class 1** means the stock should be bought, as its price will experience a significant uptrend after the announcement of financial statements.

- **Class 0** means no action should be taken, as its price will stay in a sideways trend or it will experience a small but risky (high volatility) uptrend.
- **Class -1** means the stock should be sold, as its price will experience a deteriorating performance.

These classes are obtained based on the Sharpe ratio, a commonly used measure of investing performance [11]:

$$SR_s = \frac{\bar{r}_s - r_f}{\sigma_s} \quad (1)$$

where \bar{r}_s is a mean return of stock s , r_f is a risk-free return, and σ_s is a standard deviation of excess returns for s .

SR_s represents the standardized excess return of an investment. Taking into calculation σ_s , it penalizes riskier investments. For the purpose of this challenge, the Sharpe ratio was calculated based on stock price movement for a period following the announcement of financial statements announcement until the end of that quarter (approximately two months). The class distribution is given in Fig. 3.

Although the temporal component is very important in stock predicting problems, this dataset does not comprise it explicitly. Namely, the information about a company name and the timestamp are omitted. Further, all instances are then shuffled randomly to mask information about a company and times. Still, the temporal component is included through the 1-year (absolute) change for each of these indicators.

To validate the obtained results, we performed a standard 80%/20% train-test split, ensuring that the class, NA, and null distributions, are maintained in both datasets. To avoid overfitting by submitting multiple solutions to the evaluation system, participants receive information about their preliminary score based on a small fixed subset of the test records after submitting a solution. The final evaluation is conducted after the challenge concludes, using the remaining test data.

The quality of submissions was evaluated using the average error cost measure with the error cost matrix given in Table I. The misclassifications where buying (class 1) is recommended instead of selling (class -1), and vice versa (class -1 instead

TABLE I: Evaluation cost matrix.

	-1	0	1
-1	0	1	2
0	1	0	1
1	2	1	0

of class 1), are penalized twice as much as misclassifications that resulted in taking no action (class 0). The rationale behind this was to penalize not just the actual loss because of the mistakenly predicted trend, but also the opportunity cost that comes from missed profit opportunity. It is also worth mentioning that using this cost matrix could be regarded as very similar to considering the challenge task as a regression problem with the MAE measure selected as the cost function.

Additionally, a special prize was awarded to the solution that achieved the highest cumulative risk-return performance:

$$CS_s = \sum PC_s \cdot SR_s \quad (2)$$

where PC_s is a predicted class for stock s .

C. The baseline solution

The baseline model was constructed to give participants a reference for the quality of their submissions. The model was trained using XGBoost [34]. The available training dataset was preprocessed to one-hot encode the categorical attribute indicating the company's sector. Two types of missing values were handled by setting them all to NA and adding new binary features to indicate their specific semantics. However, an investigation of the attribute importance for the final model did not reveal their substantial impact on the model's predictions. In the future, it could be worthwhile to extend this analysis using various methods to determine the significance of distinguishing between the considered two types of missing values [35].

The model's hyperparameters were not tuned extensively. The impact of several settings was checked, however, a notably large variance in evaluation results on small random subsets of training data was noticed. The final settings involved changing the default value of the learning rate η to 0.001, and the maximum depth of trees to 6. Moreover, strong regularization was enforced by subsampling features during the construction of trees with the factor 0.5 and setting the λ and α parameters to 10 and 100, respectively. Lastly, due to the observed high instability of predictions, the misclassification risk had to be taken into account. The predictions were adjusted by lowering the classification threshold for class 0, i.e., instead of simply selecting the class with the highest probability, the prediction was set to class 0 if its marginal probability was greater than 0.15. The final result for such predictions achieved a score of 0.8548, which gave it 21st position in the final ranking.

D. Participation statistics

The challenge attracted 194 teams comprising 259 individuals, which makes it one of the most popular in the history of FedCSIS. 77 enrolled teams were deemed active,

TABLE II: Final results. Preliminary and Final score columns show the average prediction costs obtained by top-ranked teams (the lower the better). Cumulative risk-return column (the higher the better) presents the results of the additional evaluation metric computed after the challenge's completion.

Rank	Team name	Preliminary	Final score	Cumulative risk-return	# subs
1	NxGTR	0.6584	0.7720	69.61	52
2	hieuvq	0.7376	0.8003	58.09	220
3	StockTrends	0.7228	0.8020	54.44	52
4	beamon	0.6980	0.8059	46.76	90
5	Pattern Pioneers	0.7822	0.8076	53.51	51
6	Team1	0.7525	0.8098	47.78	20
7	Stokastik Heinz	0.7970	0.8187	52.14	27
8	Data_Bombers	0.7624	0.8237	47.07	53
9	The Singleton	0.7921	0.8259	45.06	5
10	No-Name	0.7426	0.8270	41.95	39
...
15	O.W.C.A.	0.8020	0.8432	70.23	18
...

having submitted at least one solution. Analysis of the IP addresses of team leaders revealed that participants hailed from 28 different countries around the world, with the highest representation observed from Germany (58), Poland (50), Italy (41), Turkey (24), and Serbia (18). The collective efforts of the contestants culminated in nearly 3,000 submitted solutions. Table II presents the final rankings, scores, and submission counts for the top-performing teams, while Fig. 4 portrays the overall daily submission trends throughout the competition. The dynamics of daily submissions witnessed noticeable fluctuations. The most active phase occurred towards the end of May, with high participant engagement remaining consistent in the final week. On the final day, we noted the highest participant activity resulting in over 250 submissions.

V. POST-COMPETITION ANALYSIS

20 teams exceeded the baseline score, thus their solutions were further analyzed. As in the previous FedCSIS challenges, most teams followed general steps in data science project methodologies [36]: data processing, data cleaning, feature engineering, feature selection and extraction, model construction, and evaluation. Unlike in previous challenges, the utilization of feature selection and extraction did not help participants enhance results substantially. As for data preprocessing and elementary feature engineering, a wide spread of techniques were exercised (normalization and simple forms of data aggregation such as count, sum, prod, std, etc.), but without significant performance score improvement since the dataset was already standardized. Furthermore, top teams employed feature extraction techniques such as PCA, correlation-based filtering, feature importance estimation based on RF, XGBoost, and wrapper-based feature selection algorithms.




Fig. 4: Daily submissions over the course of the competition.

Most top teams distinguished null and NA missing values using one-hot encoding, setting different tags, or adding dummy columns. Missing values were regularly treated using standard techniques such as imputation with median, mean, zero, 2NN, or omitted if there were too many missing values in the column. Teams treated the challenge problem as classification, regression, or a combination of these two tasks. On the one hand, several teams were applying traditional ML models, e.g. MLP (in 3 solutions) or SVM (1 solution), while one team had tested a novel Kolmogorov-Arnold network. On the other, prevailing algorithms in the best solutions are more advanced, e.g., XGBoost (6 solutions), Gradient Boosting Machines (GBM, 4 solutions), LightGBM (4 solutions), and RF (3 solutions). Notably, several winning teams achieved solid performance scores using the AutoML approach. Despite the rising trend of DL usage in stock trend prediction in the research community, a low number of teams have submitted DL-based solutions. Still, there was an attempt to solve the problem with LSTM, as expected. However, as for this challenge, DL models underperformed compared to EL algorithms.

This year's challenge is unique compared to the previous ones [10], as it included two different measures of success (1, 2). The first one was used to determine the most successful predictors and to define the final rankings. The second was used to determine the most successful investor among participants. One could think that being the most successful predictor would inevitably lead to the best investment results. However, this is not true. The investment result depends not only on how accurately one predicts the trend but also on the magnitude of change of hits (correctly predicted trends) and misses (incorrectly predicted trends). Therefore, it is more significant to correctly predict the trend for stocks that generate

returns of higher magnitude (both positive and negative).

Regarding the quality of the predictor, the best three teams were NxGTR, hieuvq and StockTrends. NxGTR had significantly better performance (lower cost function) than the rest of the top 10 teams. It is also worth mentioning that all teams from the top 10 had lower final scores than the preliminary ones, which can indicate an overfitting problem. Regarding the investment performance, the best team was O.W.C.A. as it achieved the highest cumulative risk-return performance. What is interesting regarding the O.W.C.A.'s result is that the team was ranked 15th in predicting stock price movements.

Table II shows a positive correlation between the quality of the predictor and its investment performance. The order of rankings for the best three predictors is the same as for their investment performance. However, in the final results, the best predictor did not achieve the best investment performance. The only plausible explanation of this result could be that the O.W.C.A.'s model was more successful in predicting stocks with more significant upward/downward movements.

VI. CONCLUSIONS AND FUTURE WORK

The focus of FedCSIS 2024 Data Science Challenge was on predicting trends (uptrend, sideways, downtrend) of stocks constituting the S&P 500 index based on fundamental financial indicators. In addition to providing an overview of the ML algorithms used in stock market prediction, this report paper includes a detailed description of the financial dataset, evaluation procedure, and the baseline model. Furthermore, we explored solutions exceeding the baseline score, including the one achieving the highest cumulative risk-return performance.

With 194 teams and nearly 3,000 solutions, this is one of the most successful challenges at KnowledgePit.ai. Participants employed diverse combinations of data preprocessing techniques. Detailed analysis revealed that the best solutions were mainly obtained using gradient boosting algorithms, such as XGBoost, GBM, and LightGBM. These algorithms outperformed both, classical ML algorithms and the examined DL models.

Due to the high level of attention our challenge received from the ML community and its field of application being consistently a trending topic in finance, the financial dataset will be publicly available for further improvements at us.fon.bg.ac.rs/data/fedcsis2024 and KnowledgePit.ai.

VII. ACKNOWLEDGEMENTS

This work is a part of SENSEI project co-financed by EU Smart Growth Operational Programme 2014–2020 under GameINN project POIR.01.02.00-00-0184/17-00.

Also, this study is supported by University of Belgrade – Faculty of Organizational Sciences and Ministry of Science, Technological Development, and Innovation of the Republic of Serbia, institutional funding, Grant no. 200151.

REFERENCES

- [1] O. Triebe, H. Hewamalage, P. Pilyugina, N. Laptev, C. Bergmeir, and R. Rajagopal, "NeuralProphet: Explainable Forecasting at Scale," *arXiv preprint*, p. arXiv:2111.15397, 2021.

- [2] M. M. Kumbure, C. Lohrmann, P. Luukka, and J. Porras, "Machine Learning Techniques and Data for Stock Market Forecasting: A Literature Review," *Expert Systems with Applications*, vol. 197, p. 116659, 2022.
- [3] T. Kehinde, F. T. S. Chan, and S. H. Chung, "Scientometric Review and Analysis of Recent Approaches to Stock Market Forecasting: Two Decades Survey," *Expert Systems with Applications*, vol. 213, p. 119299, 2023.
- [4] M. Aché, A. Janusz, K. Zbikowski, D. Ślęzak, M. Kryszkiewicz, H. Rybinski, and P. Gawrysiak, "ISMIS 2017 Data Mining Competition: Trading Based on Recommendations," in *Foundations of Intelligent Systems – 23rd International Symposium, ISMIS 2017, Warsaw, Poland, June 26-29, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10352, 2017, pp. 697–707. [Online]. Available: https://doi.org/10.1007/978-3-319-60438-1_68
- [5] W. Jiang, "Applications of Deep Learning in Stock Market Prediction: Recent Progress," *Expert Systems with Applications*, vol. 184, p. 115537, 2021.
- [6] G. Sonkavde, D. S. Dharrao, A. M. Bongale, S. T. Deokate, D. Doreswamy, and S. K. Bhat, "Forecasting Stock Market Prices Using Machine Learning and Deep Learning Models: A Systematic Review, Performance Analysis and Discussion of Implications," *International Journal of Financial Studies*, vol. 11, no. 3, p. 94, 2023.
- [7] D. Kumar, P. K. Sarangi, and R. Verma, "A Systematic Review of Stock Market Prediction Using Machine Learning and Statistical Techniques," *Materials Today: Proceedings*, vol. 49, pp. 3187–3191, 2022.
- [8] N. Rouf, M. B. Malik, T. Arif, S. Sharma, S. Singh, S. Aich, and H.-C. Kim, "Stock Market Prediction Using Machine Learning Techniques: A Decade Survey on Methodologies, Recent Developments, and Future Directions," *Electronics*, vol. 10, no. 21, p. 2717, 2022.
- [9] I. K. Nti, A. F. Adekoya, and B. A. Weyori, "A Comprehensive Evaluation of Ensemble Learning for Stock-Market Prediction," *Journal of Big Data*, vol. 7, no. 1, p. 20, 2020.
- [10] A. Janusz and D. Ślęzak, "KnowledgePit Meets BrightBox: A Step Toward Insightful Investigation of the Results of Data Science Competitions," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 393–398. [Online]. Available: <https://doi.org/10.15439/2022F309>
- [11] K. Olorunnimbe and H. Viktor, "Deep Learning in the Stock Market – A Systematic Survey of Practice, Backtesting, and Applications," *Artificial Intelligence Review*, vol. 56, no. 3, pp. 2057–2109, 2023. [Online]. Available: <https://link.springer.com/article/10.1007/s10462-022-10226-0>
- [12] P. Balasubramanian, C. P. S. Badarudeen, and H. Sriraman, "A Systematic Literature Survey on Recent Trends in Stock Market Prediction," *PeerJ Computer Science*, vol. 10, p. e1700, 2024.
- [13] H. H. Htun, M. Biehl, and N. Petkov, "Survey of Feature Selection and Extraction Techniques for Stock Market Prediction," *Financial Innovation*, vol. 9, no. 1, p. 26, 2023.
- [14] A. Janusz, A. Krasuski, S. Stawicki, M. Rosiak, D. Ślęzak, and H. S. Nguyen, "Key Risk Factors for Polish State Fire Service: A Data Mining Competition at Knowledge Pit," in *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems, Warsaw, Poland, September 7-10, 2014*, ser. Annals of Computer Science and Information Systems, vol. 2, 2014, pp. 345–354. [Online]. Available: <https://doi.org/10.15439/2014F507>
- [15] E. Zdravetski, P. Lameski, A. Kulakov, and D. Gjorgjevikj, "Feature Selection and Allocation to Diverse Subsets for Multi-label Learning Problems with Large Datasets," in *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems, Warsaw, Poland, September 7-10, 2014*, ser. Annals of Computer Science and Information Systems, vol. 2, 2014, pp. 387–394. [Online]. Available: <https://doi.org/10.15439/2014F500>
- [16] J. Lasek and M. Gagolewski, "The Winning Solution to the AIAA'15 Data Mining Competition: Tagging Firefighter Activities at a Fire Scene," in *2015 Federated Conference on Computer Science and Information Systems, FedCSIS 2015, Łódź, Poland, September 13-16, 2015*, ser. Annals of Computer Science and Information Systems, vol. 5, 2015, pp. 375–380. [Online]. Available: <https://doi.org/10.15439/2015F418>
- [17] M. Grzegorowski and S. Stawicki, "Window-based Feature Extraction Framework for Multi-sensor Data: A Posture Recognition Case Study," in *2015 Federated Conference on Computer Science and Information Systems, FedCSIS 2015, Łódź, Poland, September 13-16, 2015*, ser. Annals of Computer Science and Information Systems, vol. 5, 2015, pp. 397–405. [Online]. Available: <https://doi.org/10.15439/2015F425>
- [18] A. Janusz, D. Ślęzak, M. Sikora, and Ł. Wróbel, "Predicting Dangerous Seismic Events: AIAA'16 Data Mining Challenge," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 205–211. [Online]. Available: <https://doi.org/10.15439/2016F560>
- [19] M. Grzegorowski, "Massively Parallel Feature Extraction Framework Application in Predicting Dangerous Seismic Events," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 225–229. [Online]. Available: <https://doi.org/10.15439/2016F90>
- [20] Ł. Grad, "Helping AI to Play Hearthstone Using Neural Networks," in *Proceedings of the 2017 Federated Conference on Computer Science and Information Systems, FedCSIS 2017, Prague, Czech Republic, September 3-6, 2017*, ser. Annals of Computer Science and Information Systems, vol. 11, 2017, pp. 131–134. [Online]. Available: <https://doi.org/10.15439/2017F561>
- [21] Q. H. Vu, D. Ruta, A. Ruta, and L. Cen, "Predicting Win-rates of Hearthstone Decks: Models and Features that Won AIAA'2018 Data Mining Challenge," in *Proceedings of the 2018 Federated Conference on Computer Science and Information Systems, FedCSIS 2018, Poznań, Poland, September 9-12, 2018*, ser. Annals of Computer Science and Information Systems, vol. 15, 2018, pp. 197–200. [Online]. Available: <https://doi.org/10.15439/2018F363>
- [22] A. Janusz, Ł. Grad, and M. Grzegorowski, "Clash Royale Challenge: How to Select Training Decks for Win-rate Prediction," in *Proceedings of the 2019 Federated Conference on Computer Science and Information Systems, FedCSIS 2019, Leipzig, Germany, September 1-4, 2019*, ser. Annals of Computer Science and Information Systems, vol. 18, 2019, pp. 3–6. [Online]. Available: <https://doi.org/10.15439/2019F365>
- [23] A. Janusz, M. Przyborowski, P. Biczuk, and D. Ślęzak, "Network Device Workload Prediction: A Data Mining Challenge at Knowledge Pit," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*, ser. Annals of Computer Science and Information Systems, vol. 21, 2020, pp. 77–80. [Online]. Available: <https://doi.org/10.15439/2020F159>
- [24] D. Ruta, L. Cen, and Q. H. Vu, "Deep bi-directional lstm networks for device workload forecasting," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*, ser. Annals of Computer Science and Information Systems, vol. 21, 2020, pp. 115–118. [Online]. Available: <https://doi.org/10.15439/2020F213>
- [25] A. Janusz, A. Jamiolkowski, and M. Okulewicz, "Predicting the Costs of Forwarding Contracts: Analysis of Data Mining Competition Results," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 399–402. [Online]. Available: <https://doi.org/10.15439/2022F303>
- [26] E. Kannout, M. Grodzki, and M. Grzegorowski, "Considering Various Aspects of Models' Quality in the ML Pipeline – Application in the Logistics Sector," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 403–412. [Online]. Available: <https://doi.org/10.15439/2022F296>
- [27] M. Czerwiński, M. Michalak, P. Biczuk, B. Adamczyk, D. Iwanicki, I. Kostorz, M. Brzeczek, A. Janusz, M. Hermansa, Ł. Wawrowski, and A. Kozłowski, "Cybersecurity Threat Detection in the Behavior of IoT Devices: Analysis of Data Mining Competition Results," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of Computer Science and Information Systems, vol. 35, 2023, pp. 1289–1293. [Online]. Available: <https://doi.org/10.15439/2023F3089>
- [28] C. Lin, "Tackling Variable-length Sequences with High-cardinality Features in Cyber-attack Detection," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of

- Computer Science and Information Systems, vol. 35, 2023, pp. 1295–1299. [Online]. Available: <https://doi.org/10.15439/2023F2385>
- [29] M. Boullé, “Predicting Dangerous Seismic Events in Coal Mines under Distribution Drift,” in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 221–224. [Online]. Available: <https://doi.org/10.15439/2016F21>
- [30] J. K. Milczek, R. Bogucki, J. Lasek, and M. Tadeusiak, “Early Warning System for Seismic Events in Coal Mines Using Machine Learning,” in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 213–220. [Online]. Available: <https://doi.org/10.15439/2016F420>
- [31] M. Trajanoska, P. Gjorgovski, and E. Zdravevski, “Application of Diversified Ensemble Learning in Real-life Business Problems: The Case of Predicting Costs of Forwarding Contracts,” in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 437–446. [Online]. Available: <https://doi.org/10.15439/2022F297>
- [32] A. Janusz, D. Kałuża, A. Chańczyńska-Krasowska, B. Konarski, J. Holland, and D. Ślęzak, “IEEE BigData 2019 Cup: Suspicious Network Event Recognition,” in *2019 IEEE International Conference on Big Data (IEEE BigData), Los Angeles, CA, USA, December 9-12, 2019*, 2019, pp. 5881–5887. [Online]. Available: <https://doi.org/10.1109/BigData47090.2019.9005668>
- [33] A. Janusz, G. Hao, D. Kałuża, T. Li, R. Wojciechowski, and D. Ślęzak, “Predicting Escalations in Customer Support: Analysis of Data Mining Challenge Results,” in *2020 IEEE International Conference on Big Data (IEEE BigData 2020), Atlanta, GA, USA, December 10-13, 2020*, 2020, pp. 5519–5526. [Online]. Available: <https://doi.org/10.1109/BigData50022.2020.9378024>
- [34] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2016, 2016*, pp. 785–794. [Online]. Available: <https://doi.org/10.1145/2939672.2939785>
- [35] A. Janusz, D. Ślęzak, S. Stawicki, and K. Stencel, “A Practical Study of Methods for Deriving Insightful Attribute Importance Rankings Using Decision Bireducts,” *Information Sciences*, vol. 645, p. 119354, 2023. [Online]. Available: <https://doi.org/10.1016/j.ins.2023.119354>
- [36] C. Schröer, F. Kruse, and J. M. Gómez, “A Systematic Literature Review on Applying CRISP-DM Process Model,” in *CENTERIS 2020 – International Conference on ENTERprise Information Systems / ProjMAN 2020 – International Conference on Project MANagement / HCist 2020 – International Conference on Health and Social Care Information Systems and Technologies 2020*, ser. Procedia Computer Science, vol. 181, 2021, pp. 526–534. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1877050921002416>

Decoding Financial Data: Machine Learning Approach to Predict Trading Actions

Yat Chun Fung, Bekzod Amonov

TU Dortmund University

Department of Statistics, Data Science

Dortmund, Germany

yatchun.fung@tu-dortmund.de, bekozod.amonov@tu-dortmund.de

Abstract—This paper presents a study on predicting stock trends using a dataset consisting of key financial indicators from 300 S&P 500 companies over a decade. Each company is characterized by 58 financial indicators along with their 1-year changes, offering valuable insights into potential trends. The objective is to develop predictive models to accurately forecast trading actions (buy, sell, hold) based on fundamental financial data. Three machine learning models—Random Forest, CatBoost, and XGBoost classifiers—were trained, employing two distinct voting mechanisms. The first voting mechanism was utilized in the competition, while the second was developed post-competition after the test labels were released. Notably, the second model was trained solely on the training data. The results demonstrate that both voting mechanisms effectively capture trends, as reflected by the average error cost measure, evaluated using the provided error cost matrix.

I. INTRODUCTION

PREDICTING stock trends has long been a crucial aspect of financial analysis, enabling investors and traders to make informed decisions about buying, selling, or holding stocks. With the advent of advanced machine learning techniques, the ability to forecast stock movements based on fundamental financial data has significantly improved.

This research aims to enhance the accuracy of stock trend predictions by developing and training various machine learning models, including Random Forest, CatBoost, and XGBoost classifiers. Ultimately, the goal is to make informed buy, sell, and hold decisions based on financial indicators using the machine learning models developed in this research.

The composition of this study is as follows: Section 2 provides a brief description of the provided training data set. Section 3 introduces the methodologies, including data preprocessing techniques, models, and prediction techniques used in this study. Section 4 provides a descriptive analysis of the training dataset, evaluates the model designs (two voting mechanisms) and their outcomes, and examines the quality of the predictions. Finally, Section 5 offers conclusions and discusses potential directions for future research.

II. DATA SET DESCRIPTION

The training dataset consists of 8,000 instances from 300 companies, each described by 58 financial indicators and their 1-year changes. These companies are categorized into 11 sectors, as indicated by the *Group* column, which is the only

categorical feature in the dataset; all other features are numerical. The dataset includes two target variables: *Perform* and *Class*, with the primary objective being to predict *Class*. The *Perform* variable is numerical and reflects the company's stock market performance, while the *Class* variable is categorical, taking values of -1, 0, or 1, corresponding to sell, hold, or buy decisions, respectively.

There are two types of missing values in the dataset: "NA" and empty strings. "NA" indicates missing information, while empty strings represent non-applicable values. A total of 2,806 rows contain missing values.

III. METHODOLOGY

A. Data Preprocessing

1) *Categorical Data Handling*: For the only categorical feature, *Group*, one-hot encoding was applied to convert it into a numerical format. This process involved creating binary columns for each unique category within the *Group* feature, allowing the model to interpret categorical data as distinct numerical values without imposing any ordinal relationship. This encoding ensures that the categorical data is appropriately represented for analysis and model training [1].

2) *Data Imputation*: As mentioned above, there are two types of missing values ("NA" and empty strings) in the data set. For the former missing values, the MICE (Multiple Imputation by Chained Equations) method was used for imputing missing data. The process begins by initializing the missing values with a placeholder (such as the mean). Then, in an iterative manner, each variable with missing data is predicted based on the other variables in the dataset. This prediction is updated in each iteration, progressively refining the imputed values until the process converges to stable estimates. This iterative refinement helps ensure that the imputations are consistent with the underlying data structure [2].

However, instead of using all features as predictors, the top three most correlated features were selected for imputing the missing values after testing. This selective approach improves the quality of imputation by reducing the influence of less relevant variables, which ridge regression alone may not fully mitigate. By focusing on the most relevant relationships within the data, this method helps produce a more robust dataset for subsequent analysis.

For the other type of missing values, represented by empty strings, a placeholder with extremely large value was used. These values were not imputed using MICE because they are not applicable to the specific rows in which they appear. By using a placeholder, we can isolate their influence on the models' performance, which is particularly effective for tree-based methods where such distinct values can be handled appropriately without distorting the analysis.

B. Model selection

1) *Random Forest*: Random Forest is an ensemble learning method, specifically a type of bagging, that aims to improve model stability and accuracy by aggregating multiple strong learners—in this case, decision trees. Each decision tree in a Random Forest is constructed using bootstrapped samples of the data, with random subsets of features considered at each node split. This approach enhances model diversity and helps prevent overfitting through majority voting [3].

The decision trees in Random Forests determine splits based on reducing Gini impurity, which allows the model to perform inherent feature selection. This not only improves the model's predictive power but also provides valuable insights into the relative importance of each feature in the dataset. Such insights are particularly useful in high-dimensional datasets like the one at hand, which contains 117 variables, where domain knowledge alone may not clearly indicate the most important features [4].

2) *XGBoost*: XGBoost (eXtreme Gradient Boosting) is another ensemble method employed in this research, specifically a boosting technique. It offers several features that make it particularly well-suited for this dataset. Firstly, XGBoost includes built-in regularization (both L1 and L2), which helps control model complexity and prevent overfitting—issues that are common with our models in this dataset. Secondly, its ability to handle missing data and its sparsity awareness are particularly advantageous, given that approximately 35% of the rows in this dataset contain missing values. This capability is crucial in managing complex datasets with a large number of features. Additionally, XGBoost's scalability and efficiency, enabled by innovations such as a sparsity-aware tree learning algorithm, parallel and distributed computing, and out-of-core computation, make it highly effective and time-efficient for hyperparameter tuning [5].

3) *CatBoost*: CatBoost (Categorical Boosting) is another ensemble method utilized in this research, specifically designed to excel in handling categorical variables within a boosting framework. It offers several features that make it particularly well-suited for this dataset. Firstly, CatBoost implements an ordered boosting technique, which mitigates the prediction shift problem commonly encountered in traditional gradient boosting methods. This leads to more accurate and stable models, particularly important given the complexity of our dataset. Secondly, CatBoost natively handles categorical data without requiring extensive preprocessing, making it an ideal choice for datasets like ours that include categorical variable *Group* [1].

4) *MLP Classifier*: The MLP (Multi-Layer Perceptron) is a type of feedforward artificial neural network composed of multiple layers of interconnected nodes, where each node is fully connected to every node in the subsequent layer [6]. In this research, the MLP is employed for the soft voting mechanism, which is used in the second model architecture.

C. Model Voting

1) *Hard Voting*: Hard voting is an ensemble technique where multiple models vote on predicted class labels, with the majority rule determining the final prediction. Each model contributes one vote, and the label with the most votes is selected. This method helps reduce overfitting and improve generalization by leveraging the strengths of different models [7]. This technique was employed in the first model architecture used during the competition.

2) *Soft Voting*: Soft voting is an ensemble technique where the predicted probabilities from multiple models are averaged to make a final prediction. Unlike hard voting, which considers only class labels, soft voting factors in each model's confidence, often leading to more accurate decisions [7]. This method reduces variance and bias, resulting in more robust performance across datasets. In this research, the MLP's weights are used for soft voting, which is employed in the second model architecture.

IV. DATA ANALYSIS

A. Descriptive Data Analysis

1) *Data Distribution*: The dataset exhibits an imbalanced *Class* distribution, with 47% of observations classified as *Class* 1 (Buy), 39% as *Class* -1 (Sell), and only 14% as *Class* 0 (Hold). The distribution of *Perform* appears symmetric and follows a bell-shaped curve, resembling a normal distribution, with a mean of 0.0341 and a variance of 0.0215 (both rounded to three sig. fig.).

According to Figure 1, observations can be categorized into three classes based on their performance values: high values (> 0.04) align with *Class* 1 (Buy), moderate values (between -0.015 and 0.04) align with *Class* 0 (Hold), and low values (< -0.015) align with *Class* -1 (Sell). Additionally, these thresholds suggest the potential for a regression task, where predicted values from regression models can be used to determine *Class* labels with the above thresholds.

In regard to the features, the distribution of the *Group* variable is notably imbalanced. For the distribution of the numerical features, please refer to the table in the appendix. It can be observed that some of the distributions are right-skewed, indicating the presence of outliers.

2) *Variable Relationships*: The relationship between *Perform* and the categorical variable *Group* was examined using boxplots. The values of *Perform* generally display similar distributions, with comparable central tendencies (median) and variability (IQR) across different groups. Notable exceptions include G3, which exhibits a negative median, and G8, which has large outliers. Additionally, G4 and G5 have outliers on




Fig. 1. Swarm plot showing the distribution of the *Perform* variable across different classes. The plot illustrates how the *Perform* values are distributed within each class. It highlights the separation of classes based on the *Perform* variable.

both ends—large and small—while G1 has one extremely high outlier.

The relationship between the target variable *Perform* and the numerical independent variables was analyzed by calculating the correlation coefficients. The highest absolute correlation was observed with variable I9 (Cash Flow from Operations to Total Assets), which had a value of 0.0762. These low correlation values indicate that there is essentially no significant linear relationship between the target variable *Perform* and the independent variables.

B. Modeling

1) *Data Preprocessing*: The same preprocessing steps were applied throughout the pipeline. The categorical variable was one-hot encoded into 11 binary columns, except in the case of CatBoost, which natively handles categorical variables. For numerical features, missing values (NAs) were first imputed using MICE, followed by standardization. Lastly, placeholders were added for other types of missing values (e.g., empty strings), with this step performed last to prevent any impact on the standardization process.

2) *Custom Prediction Rule*: A custom prediction rule was applied to derive the labels from the probability vector after voting. If the probabilities for both -1 and 1 are below 0.5, the model predicts 0 to minimize expected loss; otherwise, it selects -1 or 1 based on the higher probability. The derivation of this rule is illustrated in the equation below.

$$\text{Expected Error} = \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{pmatrix} = \begin{pmatrix} \pi_2 + 2\pi_3 \\ \pi_1 + \pi_3 \\ 2\pi_1 + \pi_2 \end{pmatrix}$$

$$\text{argmin}_{-1,0,1} \text{Expected Error} \Rightarrow \begin{cases} -1 & \text{if } \pi_1 \geq 0.5 \\ 0 & \text{otherwise} \\ 1 & \text{if } \pi_3 \geq 0.5 \end{cases}$$

where π_i stands for the probability for Class i , $i = \{-1, 0, 1\}$.

3) *Model Architecture*: For the first model, an ensemble learning approach was implemented by fitting three classification models: Random Forest, XGBoost, and CatBoost. First, the predicted labels were adjusted using custom prediction rules, followed by a voting process. The voting rules are as follows: If the predictions of all models coincide, any of the predictions may be chosen. In cases where all predictions differ, the prediction from the Random Forest classifier, which demonstrated the best preliminary results (only after the custom prediction rules), is selected. When the prediction from the strongest model (Random Forest) coincides with that of any weaker model (XGBoost and CatBoost), that prediction is chosen. If the predictions of the weaker models coincide but differ from the strongest model, the consensus of the weaker models is selected.

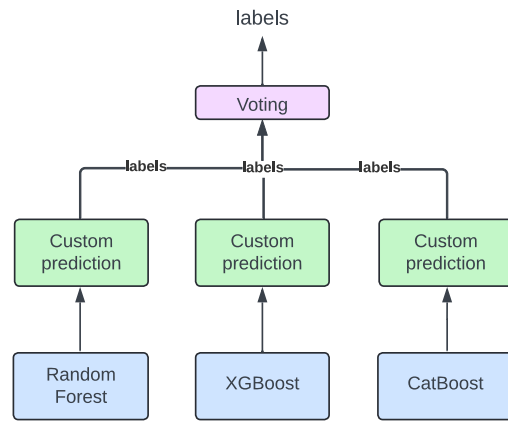


Fig. 2. Voting mechanism where predictions are adjusted based on custom rules. The Random Forest prediction is favored if all predictions differ, while the consensus of weaker models is chosen when they align but differ from the strongest model

For the second model, an ensemble learning approach by fitting three classification models was implemented: Random Forest, XGBoost, and CatBoost. Each model produces a probability vector, representing the predicted probabilities for each class. These three probability vectors are then used as inputs to a MLP classifier. The MLP has a single hidden layer consisting of 20 neurons, with the output layer corresponding to the true class labels.

After training the MLP, we extract the weight matrix associated with the connections between the input layer (9 nodes corresponding to the probabilities from the three models) and the hidden layer (20 neurons). This weight matrix is of size 9x20. To derive feature importance, we sum the weights row-wise across the matrix, resulting in a 9x1 column vector.

This column vector is then used as a set of weights to perform weighted soft voting, where each probability vector from the three models is multiplied by its corresponding weight. The

final predicted class label is determined by applying custom prediction rules to the aggregated weighted probabilities.

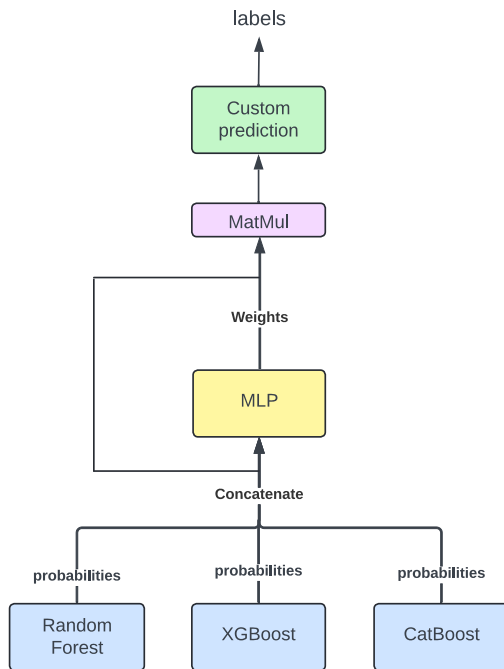


Fig. 3. The model architecture combines probability vectors from Random Forest, XGBoost, and CatBoost models, which are input into an MLP with one hidden layer. The extracted weight matrix is used for weighted soft voting to determine the final class label.

C. Results

1) *Model Evaluation and Error Diagnosis*: Among the three models (Random Forest, XGBoost, and CatBoost), XGBoost demonstrates the best overall performance, achieving the highest accuracy, weighted precision, weighted recall, and F1-score. It also has the lowest error (both preliminary and final) according to the cost matrix (0.8218 and 0.8355). Specifically for Class -1, XGBoost has the second-highest precision but the highest recall, indicating its strong ability to identify Class -1 observations. However, the lower precision compared to CatBoost suggests a potential issue with overfitting. Despite this, XGBoost achieves the highest F1-score for Class -1, balancing its recall and precision effectively.

For Class 0, XGBoost achieves the highest recall and precision among the models. This can be attributed to the fact that it is the only model that makes Class 0 predictions (thanks to the boosting algorithm probably), producing 8 predictions in total. However, only 2 of these predictions are correct, resulting in a precision of 0.25. The F1-score of 0.01 reflects XGBoost's challenges in both identifying Class 0 data and accurately predicting it. Despite these difficulties, XGBoost still has the highest F1-score for Class 0 compared to the other

models. The poor performance for Class 0 can be explained by the limited number of observations (1136 out of 8000, or 14.2%), which provides insufficient data for the model to effectively learn the patterns in the predictors. For Class 1, XGBoost has the lowest recall but the highest precision among the models. Overall, it ties with Random Forest for the second-highest F1-score. Compared to Class 0 and Class -1, the recall, precision, and F1-score for Class 1 are significantly higher across all three models. This can be attributed to the abundance of Class 1 observations in the training data (3768 out of 8000, or 47.1%), which provides ample instances for the models to effectively learn the patterns associated with this class.

For CatBoost, all performance metrics except the weighted F1-score rank second best. CatBoost excels in identifying and predicting Class 1, as indicated by its highest F1-score among the three models as shown in Table I. However, it struggles with Class 0, as evidenced by a 0 F1-score, despite assigning more importance to Class 0 during training. This may be due to the majority classes having stronger signals, making their patterns easier to learn. As CatBoost's boosting algorithm focuses on minimizing overall error, it might prioritize the majority classes, leading to fewer correct predictions for the minority class. Additionally, CatBoost performs the worst for Class -1, as shown by its lowest F1-score for that class. Its final result is 0.8465, which is the second best among the three models.

For Random Forest, the F1 scores for Classes -1, 0, and 1 all rank second, indicating a balanced performance across the classes. Despite this apparent balance, Random Forest has the lowest accuracy, weighted precision and recall, with a final result of 0.8575, the worst among the models. A notable issue is the model's failure to predict any instances of Class 0, which makes up 14% of the dataset. This problem likely arises due to the nature of the data, which may lack strong, consistent signals in the predictors. As a result, the decision trees within the Random Forest may produce very different results, causing the majority voting process to fail.

The bagging algorithm used in Random Forest could further exacerbate this issue. By training on different subsets of data, the model may inadvertently reduce the representation of Class 0 even more, leading to its omission in the final predictions. Additionally, if the predictors for Class 0 are weak or overlap significantly with other classes, the trees may not learn to split on features that identify Class 0 effectively.

Our first final model achieved a score of 0.8059, placing 4th on the leaderboard. Compared to the three individual component models, the final model has lower accuracy, weighted recall, and F1-score. However, it demonstrates a higher weighted precision, which contributes to the improved score. This precision boost is also influenced by our custom prediction rules, where misclassifying Class 0 results in a lower error (only 1), making it less costly.

The custom prediction rule dictates that the model will predict Class 0 unless the confidence (probability) for either Class -1 or Class 1 is high enough (≥ 0.5). This approach

significantly increased the number of Class 0 predictions while reducing the predictions for Classes -1 and 1 (148 predictions for Class -1, 1204 for Class 0, and 648 for Class 1).

For Class 0, the precision is lower than that of XGBoost. This is because the custom prediction rule leads to many Class 0 predictions even when the model is not confident. Consequently, many of these predictions are made not because the model is certain about Class 0, but because the probabilities for Classes -1 and 1 are not sufficiently high. This uncertainty reduces the precision for Class 0. On the other hand, the recall for Class 0 increases significantly due to the higher number of Class 0 predictions.

For Class 1, the first final model achieves better precision than XGBoost (which had the highest precision for Class 1 among the individual models). This improvement is because the number of Class 1 predictions decreases, but the correct predictions for Class 1 do not decrease as much, thanks to hard voting. This results in a higher precision for Class 1.

Regarding Class -1, the precision remains the same as CatBoost (the model with the highest precision for Class -1). This is because the decrease in correct predictions for Class -1 is more pronounced than the overall decrease in the number of Class -1 predictions. This outcome can be attributed to the limited number of Class -1 instances in the training data. Even when the model correctly identifies Class -1, the confidence often falls below 0.5, leading to these predictions being overridden by the custom rule. As a result, the precision for Class -1 does not improve.

For the second final model, which was developed after the competition, the final result was 0.797. The model exhibits higher or equal precision across all three classes compared to the first final model, thanks to the soft voting approach that takes into account the probabilities from each component model when making decisions. Specifically, the model made 168 predictions for Class -1, 1,415 for Class 0, and 417 for Class 1.

Compared to the first model, the second model produces even more Class 0 predictions, likely to minimize the error when the model is uncertain. It also makes more correct predictions for Class -1, resulting in increased precision and recall for that class. However, Model 2 makes fewer predictions for Class 1. Most of the predictions that were reduced were originally misclassified, leading to an increase in precision but a decrease in recall for Class 1. The overall improvement in the score is likely due to the increased precision for Classes -1 and 1, as well as the increased number of Class 0 predictions, reflecting the model's cautious approach when uncertain.

V. CONCLUSION

This study investigated stock trend prediction using key financial indicators from 300 S&P 500 companies. Three machine learning models—Random Forest, CatBoost, and XGBoost—were employed with two distinct voting mechanisms. While XGBoost delivered the best overall performance, our custom Model 1 and Model 2 achieved better final results by

TABLE I
PERFORMANCE COMPARISON OF DIFFERENT MODELS ACROSS VARIOUS (WEIGHTED) METRICS

	Acc.	Precision	Recall	F1-Score	Prelim.	Final
Baseline	0.1425	0.0203	0.1425	0.0355	0.8564	0.8575
RF	0.5000	0.4244	0.5000	0.4458	0.8465	0.8575
XGBoost	0.5100	0.4713	0.5100	0.4659	0.8218	0.8355
CatBoost	0.5055	0.4381	0.5055	0.4317	0.8564	0.8465
Final Model 1	0.3155	0.4948	0.3155	0.3211	0.6980	0.8059
Final Model 2	0.2810	0.5345	0.2810	0.2923	0.7673	0.7970

Remark: Acc. denotes accuracy, Prelim. refers to the preliminary score on the leaderboard, and Final represents the final score on the leaderboard

TABLE II
COMPARATIVE F1-SCORES BY CLASS FOR DIFFERENT MODELS

	Baseline	RF	XGBoost	CatBoost	Final Model 1	Final Model 2
-1	0.00	0.41	0.46	0.34	0.17	0.21
0	0.25	0.00	0.01	0.00	0.24	0.25
1	0.00	0.61	0.61	0.63	0.47	0.38

TABLE III
COMPRATIVE PRECISION SCORES BY CLASS FOR DIFFERENT MODELS

	Baseline	RF	XGBoost	CatBoost	Final Model 1	Final Model 2
-1	0.00	0.48	0.49	0.52	0.52	0.58
0	0.14	0.00	0.25	0.00	0.15	0.15
1	0.00	0.51	0.52	0.50	0.58	0.61

TABLE IV
COMPARATIVE RECALL SCORES BY CLASS FOR DIFFERENT MODELS

	Baseline	RF	XGBoost	CatBoost	Final Model 1	Final Model 2
-1	0.00	0.35	0.44	0.26	0.10	0.13
0	1.00	0.00	0.01	0.00	0.63	0.73
1	0.00	0.77	0.72	0.86	0.40	0.27

effectively managing prediction uncertainties. Model 2, developed post-competition, demonstrated further improvements in precision across all classes, underscoring the effectiveness of soft voting. Despite the challenges posed by class imbalance, particularly for Class 0 and -1, our approach successfully captured significant trends, as reflected in the final scores. Looking forward, a promising direction would be to explore stacking as an ensemble method, using the probability vectors from the three models as input to an MLP classifier to produce a refined probability vector, followed by custom prediction rules to determine the final labels.

REFERENCES

- [1] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "Catboost: Unbiased boosting with categorical features." *Advances in Neural Information Processing Systems*, vol. 31, pp. 6638–6648, 2018.

- [2] S. Van Buuren and K. Groothuis-Oudshoorn, "Mice: Multivariate imputation by chained equations in r," *Journal of statistical software*, vol. 45, no. 3, 2011.
- [3] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [4] A. Liaw and M. Wiener, "Classification and regression by randomforest," *R news*, vol. 2, no. 3, pp. 18–22, 2002.
- [5] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016, pp. 785–794.
- [6] K. Gurney, *An introduction to neural networks*. CRC press, 1997.
- [7] L. I. Kuncheva, *Combining Pattern Classifiers: Methods and Algorithms*. John Wiley & Sons, 2004.

Searching Stable Solutions For Stock Predictions: A Stacking Approach

Ty Gross, Arthur Allebrandt Werlang, Apeksha Poudel, Julian Roß
Technische Universität Dortmund
Department of Computer Science
Dortmund, Germany

{ty.gross, arthur.werlang, apeksha.poudel, julian.ross}@tu-dortmund.de

Abstract—The goal of the competition is to predict stock positions for holding, selling or buying stocks of companies from the S&P 500. Firstly the data is read in and the missing values are imputed with the median. Categorical data is one-hot encoded. A classification approach with mainly tree based methods is used. The models used are HistGradientBoosting, XGBoost, MLP and SVC whose parameters are chosen and modified through a grid search. For the stacking the models’ prediction results are summed up and the result is mapped to the three positions. It is found that the result is a bit overfitted to the competition’s test data which makes sense in regard to it being a competition. The stacking improves the score drastically. Concluding it can be said that machine learning models can hint in the right direction when it comes to handling stocks but fail at giving good financial advice.

I. INTRODUCTION

THIS paper concludes the results of team "TUmany Data" in this year’s Knowledge Pit data mining challenge titled "FedCSIS 2024 Data Science Challenge: Predicting Stock Trends".¹ A given dataset of the S&P 500 index consisting of financial indicators of 300 companies for the last 10 years has to be analysed. The results are predictions of buy, sell and hold positions. Firstly the preprocessing of the data has to be done. This includes the reading in, imputation of missing values and the encoding of the dataset. A median imputation method is used and the values are one-hot encoded.

In the second step, the data is analysed and the predictions are evaluated. Different classification and regression models are tested but it is noticed that the performance of classification boosting and classification tree models are better than the regression models. So based on their performance, only 4 of the models i.e HistgradientBoostingclassifier, XGBoost, MLP Regressor and SVC are finalized for data imputation and evaluation. The parameters for those models are tuned using Grid Search algorithm to find the optimal parameters. The models are then evaluated by first using the cross validation and then the error function provided by the competition.

In the third step, the results are stacked to refine the predictions and to better the score. For the stacking the sum

of the different models is taken and the result is mapped to the three decisions of buying, selling and holding the stock. Because of the given penalty for wrongly predicting a stock, the mapping is skewed towards the holding position because it has the overall lowest cost.

II. METHODS

In this paper, Python is primarily used as the programming language of choice due to its wide range of library support. The Pandas and Numpy libraries [2] are used to store and process the data and are commonly used in data science. Additionally, the Scikit-learn library [3] provided many of the model implementations used in this work. During initial experimentation, the MissForest imputation library is also utilized with minor changes to fix function names that have changed in its dependencies. However, this method proved to not be effective and is not included in the final prediction models (see III-B) [4].

III. PREPROCESSING

Before the dataset can be evaluated, it must first undergo a preprocessing step to transform the data into a format that can be evaluated easily. This section highlights the properties of the given dataset, as well as encoding and imputation techniques performed on the dataset during the preprocessing stage.

A. Dataset

The dataset for the competition consists of an 8000-line training CSV file, a 2000-line unlabeled testing data CSV file, and dictionary datasets that have names for the column names. In its initial state, the given dataset has two types of missing values: "empty" and "NA". Each row of the data represents one financial statement. The columns contain 58 key financial indicators and a perform and class column that can only be found in the training data and not the test data. The class column contains three classes, "sell"/"hold"/"buy", labelled as "-1"/"0"/"1" accordingly.

Figure 1 gives a visual representation of the information stated above. From the graph, it is clear that some columns have a significantly higher number of NA/missing values compared to empty values. Columns such as I21, I48, I50, dI21, dI48, and dI50 exhibit particularly high counts of NA/missing

¹The challenge was held on the knowledgepit.ai platform. It was organised by the Conference on Computer Science and Intelligence Systems (FedCSIS) and sponsored by Yettel.Bank as well as the Conference on Computer Science and Intelligence Systems series. The authors did not benefit financially or in form of other endorsements from this challenge. For more information about the challenge, see [1]

values. This visualization helped to identify which columns in the dataset required more attention for data imputation due to their high number of missing entries.

The financial indicators that comprise a majority of the columns are numerical values that are read in as floating point data types. The only column that is not numerical is the group column which has categorical datatypes, representing which financial sector the statement is from.

B. Data Imputation

There are multiple imputation methods for imputing missing values. In this section, a brief analysis is done to determine the best imputation method for this particular dataset. For that, a handful of imputation methods are tested against a few machine learning algorithms. The results are measured with the testing error described in the challenge (see [5] and [1]). The following imputation methods are analysed:

Imputation with mean, median, mode, random values, using a missing forest regression as well as disregarding all missing values.

The following machine learning algorithms are used to compute the testing error:

Decision Tree Regressor, Random Forest Regressor, SVR, Decision Tree MLP, Decision Tree Gradient Boosting, Bayes Ridge Regressor, Gradient Boosting, Hist Gradient Boosting, Decision Tree Methods, Ada Boost Regressor, Bagging, Gaussian Naive Bayes Negressor, SVC and an MLP.

The results of this testing matrix can be found in graphic 2. It can be seen that the tree-based methods generally perform better than the non-tree-based methods. This finding is independent of the choice of the imputation and is further built on in the evaluating process (see section IV). The decision tree gradient boosting had the best results with the imputation method mean and missing forest. When comparing the imputation methods, it can be seen that the random imputation and the mode imputation perform the worst. This is unsurprising for the random imputation. The bad result for the mode can be explained by the random-like imputation of the values. This is because most of the values occur once and are not correlated in any way with the missing values, thus being random-like. Disregarding data on a large scale generally makes the evaluation of the final model worse, and thus, this approach is not further pursued. The last three remaining imputation methods do not differ much. All three methods of regression, median, and mean, calculate similar values based on the existing values. Because of the simplicity, in the following evaluation, the median is used as an imputation method.

In addition to imputing the missing values, indicator columns are added to the dataset to indicate whether a value are imputed or not. This ensures that the model does not lose information about whether a value is missing when it is making its prediction.

C. Encoding non-numerical data

Since many models require numerical data to function, as the first step, the categorical values have to be encoded

numerically. A one-hot encoding method is used for the "Group" column, which converted the previous string values into an indicator column for each possible group. In the indicator column, if a financial statement belongs to a group, it is marked with a one in the group's column; otherwise, it is set to zero. This ensures that ML models requiring numerical data function properly on this dataset while the information about the "Group" column is maintained.

IV. EVALUATING MODELS

Once that dataset is put through the preprocessing stage, it is ready to be fed into various models. In this section, the various models and tuning methods are introduced. Then, the cross-validation method used is discussed. Finally, the evaluation metric is used to determine how well a model performs is outlined.

A. Classification vs Regression / models

Different Classification and regression models, such as Gradientboosting, Decision Tree, ADABOosting, MLP, XGBoost, Missforest, Linear Regression, etc., are applied to impute and evaluate the data. It is observed that classification tree-based models and boosting models delivered superior performance than the regression models and thus are selected to be further tuned.

Based on the performance, following models are used for further analysis:

- HistGradientBoosting Classifier [6]
- XGBoost [7]
- MLP Regressor [8]
- SVC (Support Vector Classifier) [9]

The hyperparameters of these models are tuned using the Grid-Search. Grid search is a traditional method of hyperparameters optimization, which simply makes a complete search over a given subset of the hyperparameters space of the training algorithm. In other words, the grid search algorithm is a complete brute-force and takes a too long time to execute [10].

Table I presents the parameter grid that are used and the best parameters are obtained from grid search for the above given evaluated models. This comparison highlights the specific hyperparameters and their optimal values that resulted in the best performance for each model, facilitating our further analysis and evaluation of error rates.

B. Cross-validation

For the evaluation of the models, a five-fold cross-validation is used. This means that the training data is split into five equally large test train splits. Using cross-validation helps reduce inconsistencies that can happen when data is split in an inconsistent way. For example, fitting a model and evaluating it against the outliers will have a much different outcome than fitting the model to the outliers and testing against the normal data. Thus, the variance for the testing error is reduced.




Fig. 1. Total missing values by column

TABLE I
PARAMETER METRICES AND BEST PARAMETERS FOR ALL EVALUATED MODELS.

Model	Parameter	Values Tested	Best Value
HistGradientBoostingClassifier	Learning rate	0.001, 0.01, 0.05, 0.1	0.001
	Max Depth	5, 7, 10, 20, 50	20
	Min samples leaf	5, 10, 20, 30	10
XGBoost	Learning rate	0.01, 0.1, 1	0.01
	Max Depth	2, 5, 10	5
	Number of Estimators	100, 500, 1000	100
MLP	Hidden layer sizes	(116,232,116), (116,116), (116, 232), (116)	(116,232,116)
	Activation	logistic, ReLU, tanh	logistic
	Solver	adam	adam
	Alpha	0.0001, 0.005	0.005
	Learning rate	constant, adaptive	constant
	Learning rate init	0.001, 0.005, 0.01	0.001
SVC	C	0.5, 1, 2, 5	1
	Kernel	linear, poly, rbf, sigmoid	rbf




Fig. 2. Results of the imputation analysis. Lower is better.

C. Evaluation (Calculating Error)

The same error function provided by the competition is used to evaluate the models. The error function is calculated based on the matrix in Figure 3. When the predicted class matches the expected class, zero is added to the error, predicting one off from the expected class results in one added to the error,

and predicting two off from the expected class adds two to the error. Finally, the error function divides the confusion matrix output by the number of predictions to get the average error. The decision to use the same error calculation as the competition is made to ensure that the same function is being optimized by the models as the competition’s scoring.

	-1	0	1
-1	0	1	2
0	1	0	1
1	2	1	0

Fig. 3. Confusion Matrix showing error for predicted vs. actual class value.

V. STACKING

A method of model stacking is used to optimize our further predictions. Model stacking allows for merging the predictions of multiple types of models into one final output by feeding the output of previous models into another model [11]. By stacking multiple models with different parameters and imputation methods, an incorrect prediction by one model can be mitigated by the other models being stacked. This can, however, also lead to worse predictions in the case in which one model is poorly chosen, skewing the overall results to being incorrect.

A. Choosing Models for Stacking

After evaluating multiple models, both using the training data and cross-validation, and the provided evaluation from the submissions to the competition website, different models are picked based on their performance.

Based on those factors, the best-performing models are chosen to be stacked, going from the assumption that the better-performing models would result in the best possible stack. The models that performed the best on the competition leaderboard and against the training data are the HistGradientBoosting Classifier, XGBoost, and an MLP Regressor.

B. Stacking the Models

Different stacking methods are tested - firstly the mode from the results of different models was taken, and where more than one mode is found, the prediction is set to zero. The intention behind this is to minimize the error function, as wrong hold predictions are less costly than wrong buy or sell predictions. Based on that, a new method, using the sum of many different results was devised.

It is noted, during a trial and error phase stacking the models, that a more robust method for the stacking is needed. As a result, based on the fact that the sum of n model predictions would go from $-n$ to n and could then be mapped, using the Pandas map function, so that results close to zero - meaning the models either agreed on the zero prediction or disagreed heavily - were set to zero, while results close to $-n$ or n , to -1 or 1 respectively. Since the competition score is calculated based on how far away the prediction is from the actual value (see 3), choosing zero when models disagree results in a lower penalty for an incorrect prediction. This method also considers all models equally as important, so no models are favoured, not even the better-performing ones.

In a practical example, when stacking six models, one would get results varying from -6 to 6 . From these results, any values between -2 and 2 are set to 0 , while values equal or above 3 and below -3 are set to 1 and -1 respectively. This way, it is ensured that for any buy or sell decision, at least 50% of the models agreed on that decision, which increased the accuracy of the predictions on the leaderboard substantially.

VI. CONCLUSION

While preprocessing the data, it is found that the imputation method only plays a minor role when evaluating this dataset. After a brief analysis, a median imputation is decided.

The final evaluation for the submitted predictions is a preliminary score of 0.7030 and a final score of 0.8304. The final stacking of the predictions relied on evaluating our models based on their leaderboard performance since the limited training data might not be an accurate representation of the final testing data. In hindsight, however, picking the models based on their leaderboard score might have led to the end results of an overfitting of the specific part of the test dataset that was used to calculate the error. This explains the discrepancy between the initial results and the end scores.

Further analysis should experiment with different imputation methods for missing and NA values rather than treating all of these values as the same. Additionally, the indicator columns used to mark missing/NA values should be modified to specify whether the imputed value was previously missing or NA.

For a more extensive research a regression approach with the "perform" column could be done. This would give the model more degrees of freedom to predict the score and to make a decision based on that score.

Generally, it can also be said that predicting stocks is a very hard task and can not be reliably done. Events that influence the stock market in politics and many other fields that are part of the global economy can not be foreseen. Therefore, the prediction on the basis of old data can not account for such events and changes. For a given stock, a volatile score could be assigned that captures the dependency on global events and represents its stability in changing times. For this, the dataset could be extended with additional data. For example, a score that captures the change for each stock regarding similar events to feed the machine learning algorithms the dependency for specific fields. For this reason, outside data could also be integrated into the dataset to help improve the predictions since other events outside of the financial world might have an influence on the stock's value.

ACKNOWLEDGMENT

We would like to thank Prof. Emmanuel Müller, TU Dortmund University and the Research Center Trustworthy Data Science and Security, as well as our tutors Simon Klüttermann, Michel Lang, Steffen Maletz, and Jonas Rieger, who introduced us to the challenge and supported us throughout the competition. We'd also like to thank FedCSIS for organizing the challenge and for inviting us to write this paper.

REFERENCES

- [1] A. M. Rakicevic, P. D. Milosevic, I. T. Dragovic, A. M. Poledica, M. M. Zukanovic, A. Janusz, and D. Slezak, "Predicting stock trends using common financial indicators: A summary of fedcsis 2024 data science challenge held on knowledgepit.ai platform," in *Proceedings of FedCSIS 2024*, 2024.
- [2] T. pandas development team, "pandas-dev/pandas: Pandas (latest version)," *10.5281/zenodo.3509134*, Apr. 2024.
- [3] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

- [4] Y. S. Y. Hindy, “Missforest (version 2.5.5).” <https://pypi.org/project/MissForest/>, Mar. 2024.
- [5] FedCSIS, “Data science challenge: Predicting stock trends,” 2024.
- [6] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, “Lightgbm: A highly efficient gradient boosting decision tree,” in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.
- [7] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, ACM, Aug. 2016.
- [8] G. E. Hinton, “Connectionist learning procedures,” *Artif. Intell.*, vol. 40, pp. 185–234, 1989.
- [9] J. Platt, “Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods,” *Adv. Large Margin Classif.*, vol. 10, 06 2000.
- [10] P. Liashchynskyi and P. Liashchynskyi, “Grid search, random search, genetic algorithm: a big comparison for nas,” *arXiv preprint arXiv:1912.06059*, 2019.
- [11] B. Pavlyshenko, “Using stacking approaches for machine learning models,” in *2018 IEEE Second International Conference on Data Stream Mining & Processing (DSMP)*, pp. 255–258, 2018.

Experimenting with manual and automated data mining pipelines on the FedCSIS 2024 Data Science Challenge

Luisa Buck, Marc Furier, Okan Mert Göktepe, Jusztiina Judák, Max Lautenbach, Gregor Munker
University of Mannheim
Mannheim, Germany
Email: {luisa.buck, marc.furier, okan.mert.goektepe, judak.jusztiina, max.lautenbach, gregor.muenker}
@students.uni-mannheim.de

Abstract—This paper reviews the 5th-best solution and results of the FedCSIS 2024 Data Science Challenge, which aimed to predict stock trends using financial indicators. It details the preprocessing, modelling, and tuning approaches and demonstrates, as well as the methods and techniques used to address the prediction problem effectively. Subsequently, the results of different experiments, including hyperparameter optimization on preprocessing steps and switching between different prediction targets, could be compared to manual experiments. Overall, a manually experienced model could be found to outperform hyperparameter-tuned pipelines.

I. INTRODUCTION

PUBLIC data science competitions and benchmarks support companies in deciding on today's data science methodology. They show the best-of-breed data science methods in specific fields and compare the most recent methodologies [2]. This year's FedCSIS 2024 data science challenge¹ targeted the financial markets. The challenge aimed to predict stock trends based on given financial indicators and one-year data per stock. The stock trends were predicted as three classes, which makes up a classification task. The given financial indicators are also represented in a tabular way, which opens up a variety of data mining methods.

The following paper will review our team's proceedings and results, "Pattern Pioneers", scoring 5th in the overall competition. As the team consisted of six members, various data mining methods were used to get the best scores. This paper aims to review the applied methods and share the overall experience of the application. The applied methods range from simple statistical methods like Naive Bayes to challenge-winning methods like XGBoost. It also includes the methodology of stacking as well as various preprocessing steps. Hyperparameter optimization was also utilized in the challenge but with the addition of including the whole data mining pipeline within the search space. TPOT and FLAML were included during the competition to utilize and compare automated machine-learning approaches.

The paper will be structured as follows. The first part will include a challenge review, data inspection, preprocessing,

and modelling. Afterwards, the results section will discuss how different methods were used in the given challenge. The work on the results was part of the class "Data Mining II". Therefore, the number of methods tested and the time for work were limited, limiting the time for exploitation of different approaches.

II. DATA MINING PROCESS

A. Challenge Review

The FedCSIS 2024 Data Science Challenge's topic was predicting stock trends buy, sell, and hold, making the data mining problem a classification problem. The dataset consisted of 300 S&P 500 companies, their stock trends and financial indicators collected over multiple years from the companies' financial statements. The only information given about the companies was their industry, but there was no information about the particular company. Therefore, the only data source that was legally usable was the FedCSIS 2024 dataset. The submissions to the challenge were scored on the stock trend with a type of mean absolute error, which could only be implied as the three stock trends were encoded numerically. Further information can be found within table I.

B. Data Inspection

The given data consisted of 10000 data points with the industry, 58 financial indicators, 58 1-year deltas of the financial indicators and two target columns. On the one side, the target column "Class" contained the stock trend, which was also the one in charge of the challenge leaderboard. On the other side, another column named perform was included in the dataset. This column was a risk-return performance measure. In the early part of the project, it was visible that the following rule set could discretize this risk-return measure:

This strict rule has already opened up a regression problem that could be used instead of classification. In addition, as the column "Class" was already encoded in -1, 0 and 1, the initial classification problem could also be transformed into a regression problem. On Figure 1. we can see the distribution of the instances based on the Performance attribute, colored by the actual class from the training data.

¹<https://knowledgepit.ai/fedcsis-2024-challenge/>

TABLE I
PERFORM DISCRETIZATION RULES

Perform	Class
$x < -0.015$	Sell
$-0.015 < x < 0.04$	Hold
$x > 0.04$	Buy




Fig. 1. Performance colored by Class

All in all, this resulted in a total of 117 influencing columns. The dataset consisted of 116 numeric and one categorical influencers. The categorical influencer industry consisted of eleven distinct values like energy or health care and therefore had no ordering, which makes it a nominal variable. FedCSIS has already provided a train-test data splitting. The split was an 80-20 split with no indication of shuffling or the timestamp of the collected data point. The test split was not consisting any of the performance-based columns "Class" or "Perform". Instead, this dataset was meant to define a live leaderboard. In addition, an unknown subset of the test set was used to evaluate the submission's overall performance.

Within the dataset description of FedCSIS, missing values are already outlined. There were two types of missing values in the dataset: one is non-available, and the other is non-applicable. Further information on which values remarks on which type of missing value was not conducted. When concentrating on the training dataset, only 5194 rows do not contain any missing value. When changing the axis, 75 columns contain at least one missing value. In addition, the most missing values are 1564, with a mean of 243 per column. This implies that the missing values are conducted by more than one column. Instead, they are spread across the whole dataset.

C. Preprocessing

The following various preprocessing steps will be outlined, as those were used in different approaches, leading to the result of our work. Except for the outlier removal, all preprocessing steps were used within Python 3.9 and the package scikit-learn². The outlier removal was done by using the package imblearn³. The data inspection already

leads to a couple of implications. First, as only one dataset contains every needed influencer and target column, no data transformation is needed. As mentioned, there was no legal possibility of extending the database through other data sources, such as stock price, because there was no inference on the particular company.

To evaluate the dataset at scale and not exceed 500 submissions, a private 80-20 split was done before the preprocessing. This relied on the 80-20 of FedCSIS. All shown pipelines of the later results were trained again on the full training dataset to use a maximum number of data. The missing values are an issue of the dataset that must first be resolved on the split dataset. Therefore, an imputation or deletion could be considered. As already outlined, the missing values are highly probable to be spread across the dataset; a deletion leads to a row loss of approx. 35%. In addition, this approach invokes the question of how to handle missing values in the test sets. Therefore, the imputation approach was chosen.

As already outlined, the goal of the work was to achieve results within a small given time frame; this work only contained simple imputation methods. On the one side, missing values were imputed with the mean and median of the remaining values of a column. On the other side, missing values were replaced by 0 or 999. The number 999 was intended to work as an outlier, marking a value as missing. In manual experiments kNN imputing was also tested.

In the second step of the preprocessing, the data was scaled. Therefore, a min-max scaler, the more robust standard scaler [1], and the quantile range-based robust scaler [4] were used in experiments. In this step, it is mentioned that the imputation with 999 will change the underlying distribution. Any effects of this will have to be considered during the evaluation of the results. This data scaling was done to all 116 numerical influencers. In addition, the industry encoded strings. As many data mining algorithms, especially numeric ones like linear regression or neural networks, do not support categorical values, the industry had to be encoded numerically. As this column contained nominal values, one correct way is to use a one-hot encoding, especially because there are only eleven distinct values [3].

In order to create a model later with good generalization capabilities, the outlier removal and feature selection steps were integrated after the necessary preprocessing steps. As the dataset contained 117 influencers, manual outlier detection based on outlier plots was not considered. Concerning the efficiency of the work, an isolation forest was utilized. A simple f-test based on linear regression, correlation or wrapper methods were used for the feature selection. This was wrapped into a k-best selector method. On classification approaches, additionally, sampling techniques like SMOTE were used.

²<https://scikit-learn.org/stable/>

³<https://imbalanced-learn.org/stable/>

D. Modelling

The core of the data mining pipeline remained the models and their settings. As already outlined, two types of problems could be solved within the challenge. On the one hand, there is regression and, on the other hand, a classification problem. Therefore, applicable models had to be chosen for each of the problems. Nevertheless, the models used, namely XGBoost, LightGBM, Random Forest and Support Vector Machine, supported both problems. A logistic regression was used in the first experiments, which is only capable of classification problems.

To solve the problem, three evolution steps of the pipeline modelling were done. First, manual experiments were used to create a baseline and an initial experience for the models and preprocessing on the dataset. Second, fully automated experiments were utilized to get good results efficiently. In this step, the results rely on TPOT⁴ and the package FLAML⁵. TPOT can cover the preprocessing steps so that no manual experiments with the preprocessing and TPOT are done. In contrast, the package FLAML only provides hyperparameter tooling that tests different models. The tweak in FLAML against other automated machine-learning packages is that FLAML prioritizes fast-processing hyperparameters. Therefore, FLAML fits rather well to gather good results in a time-efficient process. The third and last evolution step utilized hyperparameter tuning. All of the tuning within the work was done on a 5-fold cross-validation to maximize the generalization capabilities of the model. This was important because the dataset for the final ranking was unknown. The decision was to use Bayesian tuning to follow the idea of efficiency. Within experiments, GridSearchCV and RandomSearchCV were also performed. The underlying search spaces were chosen iteratively.

In addition to the hyperparameter of the models, all preprocessing steps were included in the tuning. Therefore, the preprocessing could be split into necessary and optional steps. The necessary steps are the scaling and encoding. For those steps, the search space consists of the presented ones in subsection II-C. The optional outlier removal and feature selection steps could be deselected within the hyperparameter tuning. When they were selected, the search space of the contamination c of the isolation forest was defined as

$$0 < c < 0.5; c \in \mathbb{Q} \quad (1)$$

and the search space of the k in k-best feature selection as

$$5 < k < 120; k \in \mathbb{Z}. \quad (2)$$

With this step, a small automated machine-learning pipeline could have been built. One minor addition was also introduced within this third evolution step: stacking the best models. This was related to good empirical performances in prior Kaggle competitions.

⁴<https://epistasislab.github.io/tpot/>

⁵<https://microsoft.github.io/FLAML/>

III. RESULTS

Throughout the project, we experimented with various models. In this section, we will discuss the three evolution steps of the pipeline modeling.

The first test was based on logistic regression, excluding NA and NaN values, which achieved an accuracy of 0.49 and an F1-Score of 0.45, with a score of 0.88 and no successful predictions for class 0. Imputing NA values with the median and deleting NaN values slightly reduced accuracy and F1-Score to 0.48 and 0.44, respectively. When both NA and NaN values were imputed with their group or column medians, accuracy improved but the F1-Score remained unchanged. Outlier detection using Isolation Forest improved accuracy to 0.50, F1-Score to 0.46, and reduced error to 0.85. Applying SMOTE for handling class imbalance reduced accuracy to 0.46 but improved error to 0.77, enabling predictions for class 0. Further experiments with feature selection based on SelectKBest and removing highly correlated features reduced the score to 0.74 and decreased accuracy and F1-Score to 0.43. The FedCSIS leaderboard showed a score of 0.84. Further techniques like PCA and SelectFromModel did not improve results.

We experimented with Support Vector Regression (SVR) for performance prediction due to its robustness to outliers and effectiveness with high-dimensional datasets. The best results were obtained with rbf kernel, Gamma set to 'scale' and C set to 0.1, resulting in a validation score of 0.89. However, when applied to the submission dataset, the score dropped to 0.79. Switching from median to K-Nearest Neighbors (KNN) imputation with two neighbors did not enhance performance, resulting in a score of 0.84. We also detected the importance of the Group attribute regarding the performance so we introduced group-based approaches. Group-based preprocessing, including mean, median, and KNN imputations, and numerical attribute scaling within groups, yielded the best result with the score of 0.80, though still lower than the initial preprocessing strategy. Despite extensive experimentation, combining OneHotEncoding, median imputation, and standard scaling, alongside an SVR model with gamma set to 'scale' and C set to 0.1, provided the best results.

For XGBoost, The final model, trained with a learning rate of 0.05, 500 estimators, and a maximum depth of 5, achieved an F1-Score of 0.61 and a validation cost of 0.5956. The submission dataset score was 0.7822. Applying SMOTE to address class imbalance and using SelectKBest for feature selection did not yield improvements, leading to the exclusion of these steps. Furthermore, we experimented with new feature generation. These new features were created based on domain-specific knowledge and combined with existing features to enhance the model's performance, but the scores worsened. Therefore, this step was also excluded.

In addition, the Random Forest approach was utilized in manual experiments. The best classification score on the FedCSIS Submission Set was 0.8861, though it did not reach the top leaderboard positions. Regarding testing and other prediction

targets, the work was focused on predicting the risk-reward performance. Switching to this target lowered the leaderboard score from 0.8861 to 0.7970, indicating high potential. That is why the following steps focussed on the regression problem of the "Perform" column.

The utilization of TPOT and FLAML seemed promising. Using FLAML with a cost-efficient approach should lower the time-to-value for best-in-class results. Nevertheless, both packages could not outperform the XGBoost classification or Random Forest regression approaches. The TPOT approach scored a best score of 0.8336 while running much longer than the FLAML package. This behaviour is reasonable, as FLAML has a specific time budget for gathering results. In a maximum of 30 minutes of training the FLAML package achieved a score of 0.8119 on the submission dataset.

In the last step and after three of four weeks in the competition, the goal was to tweak the last performance out of the models. Therefore, hyperparameter tuning, like in subsection II-D presented, was used. In addition to the scikit-learn models, FLAML was integrated as another machine-learning model. Overall, the best-performing model within this step was a random forest with a score of 0.7921 on the leaderboard. Other models like XGBoost or even the FLAML approach were behind this pipeline.

Like already outlined, we explored the use of stacking to combine the strengths of multiple predictors. Unfortunately, this idea was implemented near to the deadline of the competition and is therefore very rudimentary. Specifically, we employed the default scikit-learn StackingRegressor with the default underlying RidgeCV regression algorithm, feeding in the predictions from our best-performing LightGBM and Random Forest models from the hyperparameter tuning before, selected based on their optimal preprocessing configurations. This approach allowed us to capture more nuanced patterns and relationships in the data, leveraging the diverse strengths of each model. The resulting stacking model yielded a score of 0.7921, demonstrating the potential benefits of combining multiple predictors. Stacking has shown promise in this project, and it could be effectively used with an even broader range of models and meta-learners in future projects to enhance our predictive capabilities further.

Overall, this left three best models: a randomly experienced XGBoost Classifier, a hyperparameter-tuned pipeline based on Random Forest and a rudimentary stacking approach. Both more sophisticated approaches were experienced as outperformed slightly by the XGBoost Classifier. As the FedCSIS challenge allowed three models for the final evaluation, those three were handed in randomly to catch the best performance experienced and the models with the highest expected generalization. The final score was 0.8076, which is very close to the results experienced before. Unfortunately, the FedCSIS challenge platform is publishing which model of the three was the best performing. It is worth mentioning that good generalization was vital in getting a top 5 rank in the challenge. The best score of 0.7822 was only 42nd on the public leaderboard, with the best score being 0.5921.

IV. CONCLUSION

Through extensive experimentation with various models and preprocessing techniques, we explored various techniques to tweak the performance to the top 5 on a final leaderboard. Different preprocessing strategies, outlier detection methods, and hyperparameter tuning influenced each model's performance. Although improvements were made, some methods like group-specific preprocessing for SVR and feature selection for XGBoost did not enhance performance. In addition, the evolution steps produced better reproducible results, but they never passed the XGBoost classifier. Overall, our best results were achieved using a combination of careful preprocessing and model-specific adjustments.

REFERENCES

- [1] Christo El Morr, Manar Jammal, Hossam Ali-Hassan, and Walid El-Hallak. *Machine Learning for Practical Decision Making: A Multidisciplinary Perspective with Applications from Healthcare, Engineering and Business Analytics*. Springer International Publishing, 2022. ISBN 9783031169908. doi: 10.1007/978-3-031-16990-8. URL <http://dx.doi.org/10.1007/978-3-031-16990-8>.
- [2] Frederic Lardinois, Matthew Lynley, and John Mannes. Google is acquiring data science community kaggle. URL https://techcrunch.com/2017/03/07/google-is-acquiring-data-science-community-kaggle/?guccounter=1&guce_referrer=aHR0cHM6Ly9kZS53aWtpcGVkaWEub3JnLw&gucce_referrer_sig=AQAAAEu9gSzQHtMGz1fxcvTfrr5VG V41GmfxVdjjnmodYOzIHNhlxLXWNY7by5UshvhMOqu7rfB4Qcx05Z5fi8vMGeIVAxYorBLu--6UN1lxAG_nNgSdNy1MNv9L3m92Fxlz8kIr5YF1Kjv9z2ErFaqh3qeHzl_2_QiWylNrJMEJsK4L.
- [3] Pau Rodríguez, Miguel A. Bautista, Jordi González, and Sergio Escalera. Beyond one-hot encoding: Lower dimensional target embedding. *Image and Vision Computing*, 75:21–31, July 2018. ISSN 0262-8856. doi: 10.1016/j.imavis.2018.04.004. URL <http://dx.doi.org/10.1016/j.imavis.2018.04.004>.
- [4] Andre Ye and Zian Wang. *Modern Deep Learning for Tabular Data: Novel Approaches to Common Modeling Problems*. Apress, 2023. ISBN 9781484286920. doi: 10.1007/978-1-4842-8692-0. URL <http://dx.doi.org/10.1007/978-1-4842-8692-0>.

Key Financial Indicators Analysis and Stock Trend Forecasting Based on a Wrapper Feature Selection Method

Chang Lin

*State Key Laboratory of Information Photonics and Optical Communications
Beijing University of Posts and Telecommunications
Beijing, China
bupt.ipoc@yandex.com*

Abstract—Predicting stock price trends is a challenging puzzle. The immediate price of a stock is affected by an uncountable number of factors. Thus there is essentially no way to accurately predict short-term stock price due to dynamic, incomplete, erratic, and chaotic data. However, by analyzing key financial indicators, it is possible to gain an accurate understanding of a company's operations, make a quantitative assessment of its value, and thus make a reasonable prediction of the long-term trend of its stock price. In this FedCSIS 2024 Data Science Challenge, participants are asked to predict the trends of the stocks which are chosen from the Standard & Poor's 500 index. In this paper, we apply a wrapper feature selection method that tightly combines the steps of feature selection and model building to result in better prediction models, and provide insight into the indicators. After selecting the best set of features, we train two kinds of gradient boost machine: multi-classification model and regression model for class and risk-return performance prediction respectively. Finally a high confidence voting strategy is used to determine the kind of trading action (buy, sell, or hold). Experimental and competition results demonstrate the effectiveness of the methodology in this paper.

Index Terms—Financial Indicator, Stock Trend Prediction, Feature Selection, Gradient Boosting Decision Tree, Strategic Voting

I. INTRODUCTION

PREDICTING stock price trends is a challenging puzzle. Researchers generally agree that there are few ways to accurately predict the direction of the stock market over the next few days or weeks, but it may be possible to make price predictions for next years with meticulous study. With the rapid development of technologies such as artificial intelligence and global digitization, the prediction of the stock market has entered a technologically advanced era. Many analysts and researchers have developed various Artificial Intelligence (including Machine Learning and Natural Language Processing) based tools and techniques to predict stock price movements and help investors in proper decision-making. For example, Leippold et al. investigate 11 machine learning method's (such as ordinary least squares regression, least absolute shrinkage and selection operator, elastic net, gradient boosted regression trees, random forest etc.) predictive power in the Chinese stock market [1]. They build and analyze a comprehensive set

of return prediction factors of Chinese market and find that the most critical factors have entirely different characteristics than the US market. They also show that machine learning methods can be successfully applied to various markets with different characteristics. Wu et al. present BloombergGPT, a 50 billion parameter language model trained on a wide range of financial data, and demonstrate that their model outperforms existing models on financial tasks by significant margins [2]. However, it's worth mentioning that it took about 53 days to train the BloombergGPT at a cost of around \$3M. Basically, all deep learning based algorithms (e.g., various time series analysis methods based on Long Short-Term Memory and its variants) require extensive training on large and versatile datasets, incurring high training costs.

In the FedCSIS 2024 Data Science Challenge: Predicting Stock Trends [3], [4], participants are asked to develop a predictive model to accurately forecast stock trend movements based on the provided financial fundamental data. The selected stocks are chosen from 11 industry sectors of the Standard & Poor's 500 index, spanning 10 years. The dataset contains 117 fields (58 key financial indicators and 58 absolute changes of these indicators, and 1 industry sector) for 300 companies. For this kind of non-time-series tabular data, conventional machine learning methods (e.g., GBM-like algorithms) are well suited for modeling and analysis. In this competition, we build two Gradient Boosting Machine (GBM) models to predict whether it is a good moment to buy, sell or hold the stock, and under what circumstances the performance of investments can be maximized, respectively. We wrapped the predictive models into our proposed feature selection framework [5], [6]. Therefore, we can eliminate the influence of ineffective indicators, and find out the factors that play a key role in the predictive models. The resulting models are concise, accurate, have strong generalization capabilities, and can be well interpreted. The similarities between our work and the work of Rakićević et al. [7] are: improve the predictor's performance and provide a deeper insight into each of the indicators used for prediction.

As usual, the competitions of KnowledgePit are always well organized. The organizers carefully reviewed the competitors'

solutions, objectively assessed the novelty of their approaches and the quality of the submitted reports. This effectively prevents improper behaviors commonly found in other competition platforms, and ensures each competitor's solution stands the test of time. Moreover, for almost all the competitions that have been held, the organizers analyzed and summarized the methods submitted by the competitors and present informative papers [8], [9]. This allows the participants to identify the shortcomings of their own approaches and learn from the strengths of others. Driven by this favorable atmosphere, we are very happy to share our findings. This paper is organized as following: In this section, we introduce the background of the research. In section II we provide the analysis of the data. In section III we present a feature selection method which embedding the gradient boosting decision tree (GBDT) algorithm into a sequential floating forward and backward framework. We select different feature subsets to train two GBDT models, and use the ensemble of these models to predict stock trends. Section IV shows the experimental results and analyzes the role of each financial indicator in the trend prediction. The last section draws the conclusions and makes some recommendations.

II. DATA ANALYSIS AND PROCESSING

The available training data in this challenge contain 8,000 instances with fundamental financial data in a tabular format. Each instance in the data represents a financial statement announcement for one of the chosen 300 companies. It contains information on the company's sector, values for 58 key financial indicators, 1-year absolute change of each indicator, target class information, and risk-return performance for a period after the announcement. The target class is a single number from the set $\{1, 0, -1\}$ that indicates the predicted trading action (buy, hold, sell correspondingly) for the event. The test data which containing 2,000 instances have the same format and naming scheme as the training data but it does not contain columns 'target class' and 'risk-return performance'. The available data contain two distinct types of missing values that have different semantics. One corresponds to non-available/missing information which is marked by "NA" string and another one can be interpreted as non-applicable (there is no value) which is just an empty string.

We use target encoding method to encode the industry sector feature. For null and NA values, we simply set them to two specific numbers (e.g., -300 and -999) that are different from the other normal values, then our algorithm can handle automatically. Then we check the correlation coefficients between the financial indicators and the forecast targets (class and performance). Table I gives the 10 most correlated features' Spearman correlation coefficients. As can be seen from this table, there is no significant correlation between any financial indicators and the forecast targets. This demonstrates that stock change trends are the result of a combination of many complex factors. It is difficult to provide a comprehensive interpretation and make an accurate prediction of the stock trends based on a limited number of indicators.

TABLE I
10 MOST CORRELATED FEATURES' SPEARMAN COEFFICIENTS

Class		Performance	
Coefficient	Feature Id	Coefficient	Feature Id
0.070830	dI6	0.102539	I57
0.064429	dI7	0.081737	dI6
0.064336	I57	0.080568	I9
0.058301	I6	0.080181	dI47
0.055730	dI57	0.080143	Group
0.054116	dI47	0.079727	dI52
0.053434	dI9	0.079392	I18
0.053369	I8	0.076056	dI7
0.051238	I18	0.073477	I4
0.050009	I9	0.070641	dI57

Fig. 1 gives the 'boxplot' of indicator I57 and dI6. As can be seen in Figure 1, the data of these two features are very concentrated in the center and have a long spread on both sides. This is very similar to a normal distribution. Many indicators have similar distribution properties. It can also be found from the figure that these features have the same distribution in the training set and the test set. By comparing the distribution of each feature, we are confident that the data in the training and test sets have the same distribution pattern.

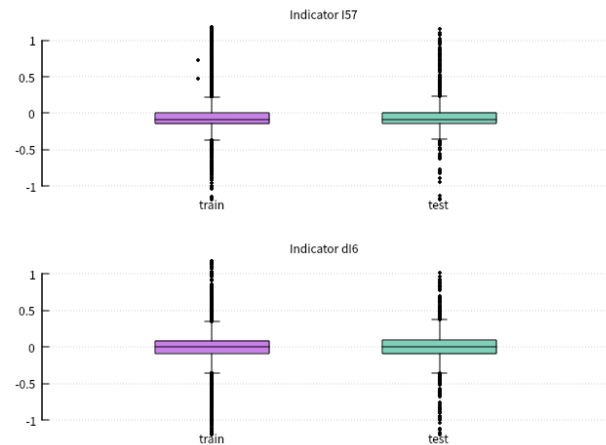


Fig. 1. 'boxplot' of indicator I57 and dI6.

Fig. 2 shows the risk-return performance of each industry sector. It is easy to find some interesting facts in this chart, for example, if one invests in energy(G3) stocks, there is a huge probability that one will lose money.

Fig. 3 shows the 'boxplot' of risk-return performance. From Fig.3 we can find that the data of risk-return performance are mainly concentrated in the range from -0.373 to 0.439, and it approximates a normal distribution. We also find that the correspondence between target 'class' and 'performance' can be described by the following equation:

$$class = \begin{cases} 1, & perform > 0.04; \\ -1, & perform < -0.015; \\ 0, & otherwise; \end{cases} \quad (1)$$




Fig. 2. risk-return performance of each industry sector.

Thus the stock trend prediction problem can be solved in two ways. First we can treat the task as a 3-classification problem for predicting a trading action (buy, sell, hold). We also can treat the task as a regression problem to fit the risk-return performance.




Fig. 3. boxplot of risk-return performance.

We also try to construct new features based on the provided financial fundamental data by adding, subtracting, multiplying and dividing. But they are found to be largely unhelpful in the prediction of stock trends.

III. METHODOLOGY

The methodology we use in this competition is very concise. We use the algorithm proposed in paper [5], [6] for feature selection. This algorithm essentially is a sequential floating forward and backward method. Its main improvement is that it embeds the GBM algorithm into the feature filtering framework. The procedure of feature selection is divided into forward and backward steps, as shown in Fig. 4.

In forward step, we sequentially select features one by one from the candidate set, add it to the selected set, use them to train the GBM, and evaluate the role played by each feature by

comparing the results of each training, then move the L best features from candidate set to selected set. L is determined by the improvement of prediction accuracy.

In backward step, we sequentially drop a feature from the selected set and use the remains to train the GBM, evaluating the role played by each feature by comparing the results of each training, then move the R worst features from selected set to candidate set or directly drop the worst features according to the evaluation scores. R is determined by the loss of prediction accuracy.

Repeat this two steps until evaluation scores cannot be improved.




Fig. 4. Sequential floating forward and backward feature selection method.

After selecting the best set of features, we train 30 GBM multi-classification models and take their average for class prediction. We also have tried quite a few other methods, none of which are superior to the GBM. And we find that usual ensemble methods do not work here, because the prerequisite of ensembling a set of weak classifiers to a strong classifier is that the accuracy of each weak classifier must be slightly greater than 50%.

Using the same procedure, we train 30 GBM regression models and take their average for performance prediction.

Finally, we use (1) to transform the 'performance' value into a classification result, which is then combined with the 'class' value by voting. We calculate the weight of each vote according to its 'performance' value or the probability of its 'class'.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

We select 50 features from the indicators and sector information to train the three-classes classifiers. The learning curve (decrease of softmax loss) of our GBM is shown in

Fig. 5. Fig. 5 also gives the training results of xgboost. Using our classifiers to predict the trading action (buy, sell, hold) yields (by 4-fold cross-validation): softmax loss = 0.982, classification accuracy = 50.7%, and average error cost = 0.8405. The hyper-parameters of our GBM and xgboost are simply set to: learning rate = 0.01, gamma = 0.01, lambda = 4.0, min_child_weight = 20, num_round = 300. In this competition, the error cost matrix is defined as:

$$\begin{bmatrix} p/t & -1 & 0 & 1 \\ -1 & 0 & 1 & 2 \\ 0 & 1 & 0 & 1 \\ 1 & 2 & 1 & 0 \end{bmatrix}$$

p : prediction, t : truth




Fig. 5. The learning curve of the classifier.

We evaluate the importance of each selected feature by computing its contribution to the total gain. Fig. 6 shows the gain contribution of 10 most important features in the classification models. As can be seen from Fig. 6, indicators such as I57(Cash Flow from Operations Pct of Capital Expenditures), I5(Excess Cash Margin), dI52(1-year Absolute Change of Cash Ratio) et al. play important role in the classification, but their importance is not decisive. The contribution of each of these 50 features to the total gain only ranged from 1% to 3.3%. We think that the most crucial thing in this procedure is that we drop a large number of invalid features which are not closely related to the classification problem and tend to degrade the performance of the classifier, thus improving the accuracy and generalization ability of the classifier.

When training the regression models to fit the risk-return performance, we chose fewer features, just 39. The learning curve (decrease of mean-square error) of our GBM is shown in Fig. 7. Fig. 7 also compares the training results with xgboost. The mean-square error (MSE) of our regression model is 0.144. Using (1) to convert 'performance' value to 'class' probability, get average error cost = 0.8015.

Fig. 8 shows the gain contribution of 10 most important features in the classification models. As can be seen from




Fig. 6. Gain contribution of 10 most important features in the classification models.




Fig. 7. The learning curve of the regression model.

Fig. 8, indicators such as I57(Cash Flow from Operations Pct of Capital Expenditures), dI58(1-year Absolute Change of Price to Cash Flow from Operations per Share), dI47(1-year Absolute Change of Cash & Cash Equivalents to Total Assets) et al. play important role in the regression, but their importance also is not decisive.

Combining the results of classification and regression, we get average error cost around 0.79x. Here we use a voting strategy that sets the prediction value to 0 by default; sets the prediction value to 1 when and only when 'performance' has a large positive value and 'class = 1' has a high probability; and sets the prediction value to -1 when and only when 'performance' has a large negative value and 'class = -1' has a high probability. This ensures that our predictions have a high degree of confidence.

We used 4-fold cross-validation in local test, and the obtained scores ranged from 0.790 to 0.799. Since the score of public leader-board was evaluated by only 200 instance, there




Fig. 8. Gain contribution of 10 most important features in the regression models.

is a significant difference between the public LB scores and the CV scores. Of our 20+ valid submissions (net of tests, obvious errors), most of our final scores largely better than 0.805, with 5 scores better than 0.790, and the best one is 0.7875.

V. CONCLUSION

The task of this competition is the prediction of stock trends. However it is more like estimating the return on investment of each company by analyzing its various financial indicators. With known data, we show that the methods proposed in this paper are concise, reliable and have excellent generalization ability. We achieved the desired results despite that we did not have enough time for fine-tuning the parameters and did not try hard to fit the test set. Our methods can provide a quantitative assessment of each financial indicator. It can be used as a good financial analysis tool.

Our research shows that the crux of stock trend forecasting is to select the indicators that are truly favorable for classification and regression in this task, and to buy or sell stocks when there has high degree of confidence, otherwise, 'hold' or just 'stay on the sidelines'. But our study also illustrates that there

are no financial indicators that can directly influence stock trends, in other words there is no obvious causal relationship between them. Stock price movements are still governed by a large number of dynamic or unknown factors. Whether the methodology of this paper can be directly applied to stock trading needs to be verified by more tests. Interested researchers are welcome to share and discuss together.

We would like to thank the sponsors and organizers for providing such valuable research data and organizing the competition with great effort.

REFERENCES

- [1] M.Leippold, Q.Wang, W.Zhou. Machine-learning in the Chinese Stock Market. *Journal of Financial Economics*, 2022, 145(2): 64-82. DOI: <https://doi.org/10.1016/j.jfineco.2021.08.017>.
- [2] S.Wu, O.Irsoy, S.Lu, V.Dabravolski, M.Dredze, S.Gehrmann, P.Kambadur, D.Rosenberg, G.Mann. BloombergGPT: A Large Language Model for Finance. arXiv:2303.17564v3 [cs.LG]. <https://doi.org/10.48550/arXiv.2303.17564>.
- [3] <https://knowledgepit.ai/fedcsis-2024-challenge/>.
- [4] Aleksandar M. Rakicevic, Pavle D. Milosevic, Ivana T. Dragovic, Ana M. Poledica, Milica M. Zukanovic, Andrzej Janusz, Dominik Slezak: "Predicting Stock Trends Using Common Financial Indicators: A Summary of FedCSIS 2024 Data Science Challenge Held on KnowledgePit.ai Platform". In: *Proceedings of FedCSIS 2024* (2024).
- [5] C. Lin. Predicting Frags in Tactic Games using Machine Learning Techniques and Intuitive Knowledge. 2023 IEEE International Conference on Multimedia and Expo Workshops (ICMEW), Brisbane, Australia, 2023, pp. 11-15, doi: <https://doi.org/10.1109/ICMEW59549.2023.00008>.
- [6] C.Lin. Tackling Variable-length Sequences with High-cardinality Features in Cyber-attack Detection. *Proceedings of the 18th Conference on Computer Science and Intelligence Systems*, Vol.35, pages 1295-1299 (2023). DOI: <https://dx.doi.org/10.15439/2023F2385>.
- [7] A.Rakićević, A.Poledica, B.Petrović. A Novel IBA-DE Hybrid Approach for Modeling Sovereign Credit Ratings. *Mathematics* 2022, 10, 2679. <https://doi.org/10.3390/math10152679>.
- [8] A.Janusz, A.Jamiołkowski, M.Okulewicz. Predicting the Costs of Forwarding Contracts: Analysis of Data Mining Competition Results. *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, ACSIS*, Vol.30, pages 399-402 (2022). DOI: <https://dx.doi.org/10.15439/2022F303>.
- [9] M.Czerwinski, M.Michalak, P.Biczuk, B.Adamczyk, D.Iwanicki, I.Kostorz, M.Brzeczek, A. Janusz, M.Hermansa, L.Wawrowski, A.Kozłowski. Cybersecurity Threat Detection in the Behavior of IoT Devices: Analysis of Data Mining Competition Results. *Proceedings of the 18th Conference on Computer Science and Intelligence Systems, ACSIS*, Vol.35, pages 1289-1293 (2023). DOI: <https://dx.doi.org/10.15439/2023F3089>.

Exploring Stability and Performance of hybrid Gradient Boosting Classification and Regression Models in Sectors Stock Trend Prediction: A Tale of Preliminary Success and Final Challenge

Ming Liu, Ling Cen, Dymitr Ruta
EBTIC, Khalifa University, UAE
{liu.ming,cen.ling,dymitr.ruta}@ku.ac.ae

Quang Hieu Vu
GREENFEED, Vietnam
hieu.vq@greenfeed.com.vn

Abstract—In the dynamic field of financial analytics, the ability to predict stock market trends is crucial for effective trading strategies, which is the task for FedCSIS 2024 Data Science Challenge: Predicting Stock Trends. This paper presents a comprehensive study on the use of hybrid gradient boosting models, incorporating both classification and regression approaches, to forecast stock trends across different sectors of the S&P 500. Utilizing a rich dataset comprising key financial indicators for 300 companies over a decade, our research aims to unravel the complexities of sector-specific trend predictions. The model leverages 58 financial indicators per company, along with their annual change metrics, to predict the future stock movements. In the preliminary phase of the competition, our hybrid model demonstrated promising results, achieving the lowest weighted error of 0.5941 among competitors. However, despite the initial success, the final phase of the model evaluation revealed a significant performance decline with the error rising above 0.84. This discrepancy highlights potential issues in model stability and preliminary performance when transitioning from a controlled to a truly unseen testing environment. This work not only underscores the complexities of predictive modeling in finance but also sets the stage for future research into creating more resilient AI-driven trading systems.

Index Terms—Stock prediction, Machine Learning, Gradient Boosting Trees, Classification, Regression, Ensemble Learning.

I. INTRODUCTION

Predicting stock market trends has been a critical challenge and a focal point of interest for investors, financial analysts, and researchers alike. The complexity and dynamic nature of financial markets make this task both intriguing and difficult. Over the decades, various traditional and computational methods [1][2] have been employed to forecast stock prices and trends, ranging from fundamental analysis of financial statements to technical analysis involving chart patterns and indicators. However, the advent of artificial intelligence (AI) and machine learning (ML) has transformed the landscape of financial forecasting, offering new insights and capabilities that were previously unattainable.

Machine learning models, unlike their traditional counterparts, can handle large volumes of unstructured data and quickly uncover complex internal patterns within them. Techniques such as regression trees [3], support vector machines

[4], neural networks [5] and long short-term memory (LSTM) networks [6] have been widely adopted to predict stock prices and trends with varying degrees of success.

Hybrid models [7]-[11] that combine multiple AI techniques or integrate machine learning with traditional financial analysis have emerged as a powerful approach to improve prediction accuracy and robustness beyond the performance of any individual even the best model.

This paper aims to explore the application of hybrid gradient boosting models, which utilize both classification and regression techniques, to predict stock trends across different sectors of the S&P 500 index to address the task given in the FedCSIS 2024 Challenge¹. We examine the preliminary success of these models in capturing the nuances of sector-specific trends and address the critical challenge of maintaining model stability and performance in diverse market conditions. With this research, we seek to contribute to the ongoing dialogue on improving the reliability and efficacy of AI-driven stock market predictions, providing valuable insights for both academic research and practical trading applications.

The structure of the remainder of this paper is outlined as follows. A concise description of the FedCSIS 2024 Challenge is provided in Section II. The analysis of data distribution and the methodologies used for feature engineering are discussed in Section III. This is followed by an explanation of the gradient boosting classification, regression and the ensemble hybrid models in Sections IV and V, respectively. The experimental results are detailed in Section VI followed by the lessons learnt from the preliminary success and the final challenge discussed in Section VII. Finally, the paper concludes with some closing thoughts and observations in Section VIII.

II. FEDCSIS 2024 CHALLENGE

The 2024 FedCSIS Data Science Challenge [12], focusing on Predicting Stock Trends, marks the 10th such event hosted by the FedCSIS Conference on Computer Science and Intelligence Systems². This special anniversary edition centers

¹<https://knowledgepit.ml/fedcsis-2024-challenge/>

²<https://fedcsis.org/>

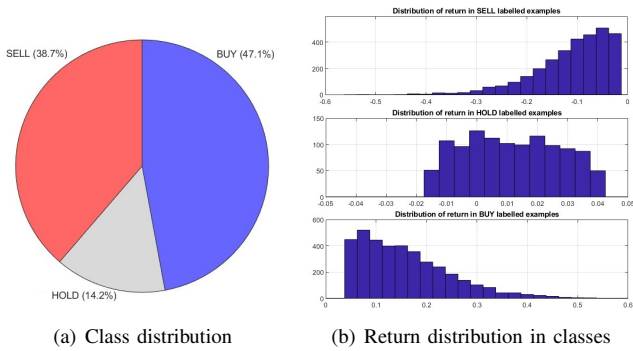


Figure 1. Distribution of trading action labels in the training dataset (8000 examples) along with the distribution of return within corresponding classes

on financial data, with participants challenged to forecast the performance of selected stocks across various industry sectors. The competition enjoys sponsorship from Yettel.Bank (former Mobi Banka)³, alongside the FedCSIS Conference itself.

III. DATA ANALYSIS AND FEATURE ENGINEERING

The challenge requires participants to devise a reliable method for predicting trading actions (buy, sell, or hold), yet with offered also a continuous return in the training set it gives a freedom of deploying classification and regression paradigm for ML model construction as long as the final outputs are crisp trading actions.

A. Training dataset

The training dataset comprises 8,000 instances, presented in a tabular CSV file format. Each data instance corresponds to a specific event—namely, the announcement of a financial statement from one of the selected 300 companies. This dataset includes information about the company’s sector, the values for 58 financial indicators and a 1-year (absolute) change for each indicator. The last two columns include the trading action (‘Class’ column), and the return performance following the announcement period (labelled as ‘Perform’ column). The distribution of classes within the training set along with the distribution of return within classes are illustrated in Figure 1.

B. Test dataset

The test dataset, which includes 2,000 instances, is formatted in the same tabular CSV file format and follows the same structure and naming conventions as the training data, however, it lacks the ‘Class’ and ‘Perform’ columns. It is important to note that not all testing set examples are necessarily in the future of all examples from the training set, which would have significantly limit the size of the testing set. However, as the organizers assured, the best care has been made to avoid temporal data leakage.

³<https://www.yettelbank.rs/en/>

C. Features

Both the training and test datasets include in total 117 features, which are divided into three categories: the Group feature that identifies one of the 11 company sectors, values for 58 critical financial indicators, and a 1-year (absolute) change for each of these indicators, as detailed in Table I.

Table I
OVERVIEW OF THE ORIGINAL FEATURES

Category	Details
Group	Financial, Industrial, Energy
	Information Technology, Consumer Staples
	Health Care, Utilities, Materials, Real Estate
	Consumer Discretionary, Communication Services
Indicators	I1, I2, ..., I58
1-Year Change	dI1, dI2, ..., dI58

D. Features aggregation and statistical features

In a search for additional features, while having limited expertise in the financial domain, we organized several raw features in groups based on their names’ similarity, as outlined in Table II, and then attempted to aggregate correlated features to achieve more stable derived features. To complete this approach we have considered several aggregation operators that have been applied to all listed groups for each instance and thereby engineered many new candidate features with potential to enhance predictability of the targets.

Table II
FEATURES AGGREGATION

Category	Aggregation
Indicators	I1-I2, I3-I4, I7-I8-I9, I29-I30-I31
	I38-I39-I40, I41-I42, I45-I46
1-Year Change	dI1-dI2, dI3-dI4, dI7-dI8-dI9, dI29-dI30-dI31
	dI38-dI39-dI40, dI41-dI42, dI45-dI46

The following list summarizes all statistical aggregators applied to group-based engineering of numerical features:

- *minimum, maximum, mean,*
- *median, sum, standard deviation*

E. Return feature

Although the return (*Perform*) column, which represents the future return of stock price movement, is only present in the training dataset, and therefore cannot be directly used for the testing set, it plays important role in designing model options and making final financial decisions. We have analyzed the return distribution per class as illustrated in Figure 1, and jointly, as shown in Figure 2 and concluded that since the return is monotonic to the ordinal encoded trading class labels (-1:sell, 0:hold, 1:buy), building a classifier on (-1,0,1) classes is equivalent to building a regression model whose outputs can be mapped back to discrete sell/hold/buy classes by simple threshold.

What is more due to the same monotonic alignment of the return and training labels as well as the fact that the model evaluation criterion uses symmetrical mis-classification cost




Figure 2. Distribution of Perform values

matrix as shown in Table III, it can be easily derived that cost-weighted mis-classification used as evaluation criterion in the competition is equivalent to the mean absolute error (MAE) of the regression model trained against discrete $(-1,0,1)$ instead of continuous targets. This discovered property gives yet another design flexibility which could be useful when building hybrid supervised ML models.

Table III
MIS-CLASSIFICATION COST MATRIX

actual \ predicted	sell(-1)	hold(0)	buy(1)
sell(-1)	0	1	2
hold(0)	1	0	1
buy(1)	2	1	0

IV. GRADIENT BOOSTING MODELS

Gradient Boosting Decision Trees (GBDT) algorithms have become a formidable and popular method in machine learning and data mining. By merging the advantages of decision trees with the technique of boosting, GBDT forms a predictive model that is both precise and robust. This methodology has been effectively utilized across several fields such as finance, healthcare, and online advertising [13][14].

For this challenge, we utilized two well-known GBDT algorithms, XGBoost and LightGBM, to develop an ensemble learning model aimed at predicting stock trends. Additionally, our team has a longstanding history of participating in data science competitions hosted by the KnowledgePit platform⁴, employing GBDT-based algorithms for tasks in classification, regression, and other areas [15] - [32], achieving remarkable success. The versatility of Gradient Boosting Decision Trees in handling various data types, along with their capabilities in feature engineering and model hyper-parameter optimization, has consistently demonstrated their effectiveness in predictive modeling across multiple fields.

To manage the task of adjusting a multitude of parameters for each specific model, we employed a rapid and efficient rotational grid search technique, an extension of the conventional

grid search method for hyper-parameter tuning [33]. Selecting the right values for hyper-parameters, including learning rate, tree depth, and regularization parameters, can markedly enhance the model's predictive accuracy and generalization capacity. And to improve the dependability of the optimal parameter configurations identified, we employed a Repeated Stratified 10-Fold cross-validation technique. This approach reduces the risk of inadvertently choosing configurations that perform exceptionally well by chance.

A. GBDT inspired Target Guided Binning (TGB)

In exploring the alternative and possibly diverse ways to build a reliable predictor we have also tested our target guided binning algorithm which has recently showed impressive predictive performance compared even to the leading gradient boosting models [16]. The simple model can be considered a combination of the Area Under the Curve (AUC) [34]-optimised 1-level singleton trees greedily merged to maximize any specific evaluation function with AUC set as a default. To attempt this model in a slightly more diverse setup we have trained it in the classification mode in two variants one using return sign as binary (buy/sell) class and another using the original buy and sell class examples only, i.e. completely excluding the hold class examples based on the rationale that the hold class examples may simply be confusing the binary classification with unstable border conditions and should therefore be trained on the strong positive (buy) and negative (sell) return examples only. We have trained both variants on all original 117 features only and achieved transformed monotonic risk features taking values from 1 (least risky - sell) to 10 (most risky - buy). We have then proceeded with constructing the model output using greedy selection of the binned features in turns maximizing the AUC at each next addition. With such selected binned features we have achieved the model returning the ordinal output of the sum of risk votes from each selected feature and the last task was to convert such output into sell, hold, buy $(-1,0,1)$ discrete labels. Since the model output is monotonic with the return we had to simply identify optimal pair of thresholds that separate the continuous domain into three bounded regions corresponding to the sell, hold and buy class. This has been achieved exhaustively yet with the fast iterative algorithm of crawling thresholds from both ends towards the middle of output range until the cost-weighted mis-classification error is minimized. Noting down these thresholds completed the TGB model build and the same thresholds have been applied to classify the testing examples. Repetitive experiments and fine-tuning of this model resulted in rather consistent rule of the best model achieved when trained without hold-class examples and allocating bottom 20% of predictions to the sell-class, 40% of top predictions to the buy-class and the remaining 40% in the middle to the hold-class. This result may somewhat come counter-intuitive to the original distribution of classes in which the hold class represents only 14% of the data, yet on balance simply reflects the cost function that penalizes double for making opposite trade mistakes and hence placing the hold

⁴<https://knowledgepit.ai/>

class as a safer bet given big uncertainty while also reflecting a buy-class bias both of which are genuinely reflected in the investment environment.

V. ENSEMBLE MODEL

In constructing the final ensemble, we utilized two core gradient boosting models: XGBoost (XGB) and LightGBM (LGBM) trained in regression mode as well as an alternative regression or classification model (initially TGB model) that we have switched on or off throughout the competition depending on the evolving leader board performance feedback. To improve the models' ability to generalize, we implemented filtering techniques. These techniques aim to diversify the classifiers by creating multiple variants, which are trained on either the full training set or specific subsets of it. These variants are then deployed on the testing set, and their outputs are collectively aggregated to form the final prediction.

To further enhance diversity and seek improved predictive performance, we trained all baseline models on different feature subsets generated by our feature engineering engine. The primary distinction between these feature subsets was that the second set included a greater number of sparse columns obtained from an extensive application of one-hot-encoding to categorical features. This approach aimed to introduce more varied and complementary information for prediction.

Additionally, to explore further opportunities for enhancing performance, we implemented an extra stacked layer of simple linear regression. This layer was trained on the outputs from the baseline models. To seamlessly incorporate this stacking layer, we split the training data into two separate segments. The initial segment was used to develop the baseline models, and the latter segment was specifically dedicated to training the parameters of the linear regression model within the stacked layer.

Ultimately, we combined the outputs from each individual model with those from the linear regression-based stacking layer by averaging them. The architecture of this final ensemble, depicted as a flow chart that illustrates the structure, is presented in Figure 3.




Figure 3. Stacking model architecture

VI. EXPERIMENTAL RESULTS

During the competition, we utilized sklearn packages, xgboost, lightgbm, and Python3 Jupyter Notebook⁵ operating on a Windows Server Virtual Machine equipped with 128G RAM and an Intel(R) Xeon(R) Gold 6230R CPU @ 2.10GHz with 2 processors for running simulations. Our approach included extensive feature aggregation and various combinations of features, where we selectively removed or filtered out specific columns as detailed in the Table IV below.

Table IV
VERSION OF FEATURE SETS

Version	Number of Features	Remarks
V1	117	Original features
V2	127	Add combined features as in Table II
V3	187	Add statistical features
V4	237	Add Top 50 importance features

The various feature sets and their respective effects on the performance of individual models, particularly within the constrained and sparse training and testing datasets, are summarized in the Table V below.

Table V
FEATURES AND MODEL PERFORMANCE

Version	LGBM	XGB	TGB
V1	0.8267	0.8218	0.8291
V2	0.7673	0.7970	0.8262
V3	0.7475	0.7624	0.8231
V4	0.7376	0.7178	0.8241

Interestingly TGB model did not seem to gain from additionally engineered features and hence was dropped from the ensemble in the subsequent hybrid model versions.

Throughout the competition, many parameter variations showed robust performance. To optimize these parameters, techniques like Grid Search was employed, which iterative explore the hyper-parameter space to identify the best combination based on cross-validated performance metrics. Our highest individual model scores were achieved using particular model settings, as detailed in the Table VI below.

Table VI
OPTIMIZED INDIVIDUAL MODEL PARAMETERS

Model	Iterations	Learning Rate	Tree Depth
LGBM	1000	0.08	3
XGB	2000	0.3	6

To accommodate the variety of company sectors indicated by the Group feature, we experimented with dividing the datasets into 11 subgroup models, applying distinct value sets to these subgroups. This approach enhanced the evaluation performance, as outlined in the Table VII below.

The final predictions were derived by averaging the results from both the stacking models and the varied individual baseline models through an ensemble technique. This method

⁵<https://jupyter.org/>

Table VII
FEATURES AND 11 SUB-GROUP MODELS PERFORMANCE

Version	LGBM	XGB
V1	0.7475	0.7624
V2	0.7079	0.7327
V3	0.6980	0.7178
V4	0.6683	0.6832

achieved a preliminary score of 0.5941, ranking as the top preliminary score among the 10% preliminary test datasets.

VII. PRELIMINARY SUCCESS AND FINAL CHALLENGE

Despite the initial success, the final evaluation phase of the model showed a decrease in performance, recording a score of only 0.841500. This decline underscores possible challenges in model stability and generalization as it moves from a controlled testing environment to a more diverse one.

After the competition organizers released the complete test datasets, we conducted a thorough evaluation of all the models. Upon analysis, we realized that the narrative of initial success followed by a subsequent drop in performance during the final phase could be attributed to two primary factors. First, there might have been an issue of overfitting, where models tuned to excel on preliminary data failed to generalize effectively to the broader dataset. Second, the variation in the test dataset's characteristics compared to the training set could have exposed weaknesses in the models' adaptability. These issues highlight the importance of robust model validation strategies and underscore the need for models that can maintain consistency across different data subsets.

To provide a more detailed explanation, the major factor concerns the sub-group models, which seemed logical given that different financial sectors may exhibit distinct financial patterns. However, a significant challenge arose due to the limited size of the preliminary test dataset, which constituted only 10% of the total data. When this data was further subdivided by Group for the subgroup models, each individual subgroup ended up with an even smaller portion of data for training. This scant amount of data likely resulted in models that were under-fitted and unstable. Such models struggle to capture the complexity and variability of their respective sectors, leading to performance issues when faced with a broader and potentially more diverse set of test data. This situation underscores the critical importance of having a sufficiently large and representative training dataset to ensure robust model training and stability.

The next concern is related to the error cost matrix⁶ as already presented in Table III, which is used as follows to compute the final error used as evaluation in the competition:

$$\text{err} = \frac{\text{confusion_matrix}(\text{preds}, \text{gt}) \cdot \text{cost_matrix}}{\text{length}(\text{gt})}$$

The other decline factor relates to the adjustments we made through post-processing techniques, which were guided by the

⁶<https://knowledgepit.ml/fedcsis-2024-challenge/>

error cost matrix. In an attempt to refine the model's performance, we utilized the cost matrix to prioritize certain types of errors over others, aligning the model's output with specific financial implications associated with different types of prediction errors. This strategy involved adjusting the model's predictions to minimize the financial risk as quantified by the cost matrix. While this approach can effectively optimize the model for scenarios represented within the training data, it may inadvertently lead to a lack of generalization on new data sets if the error characteristics differ. This reliance on post-processing based on the cost matrix can potentially skew the model's ability to predict accurately in diverse real-world scenarios, as the adjustments might not align well with the actual error distribution in unseen data.

Accordingly, we undertook some adjustments to the distribution of the labels in our dataset. Although for the TGB model we have observed the optimal predicted class distribution to be (20%,40%,40%), we have observed that for gradient boosting models, and the whole hybrid design, squashing the prevalence of hold-0 while elevating the remaining buy and sell classes seems to elevate the performance. This maneuver originating from the model fine-tuning seemed somehow to fit the preliminary set really well as we have achieved unprecedented gains in the preliminary set evaluations clearly capturing the buy and sell class examples well from that limited set. In the competitive conditions when the thorough cross-validation testing can be costly and time consuming to vet such quick new post-processing discoveries, such last minute leader board score following could lead to a classic overfitting trap that could quickly compromise generalization ability of otherwise really good ML model. Trusting the representative nature of the preliminary set we have followed through with the model adjustments that have placed our score well ahead of competition and clearly must have optimized for specific characteristics of the preliminary data, at the cost of the model's ability to perform consistently across a more comprehensive dataset. One could wonder, regardless, how is it possible to achieve such a high score for the preliminary set that was still hidden for the participants? A possible explanation could be that our complex hybrid model with several layers of grouping may have discovered data snooping leads that connected data points of the same stock from different moments in time. Whatever the reason, consequently these alterations contributed to the preliminary success yet ultimate demise of our model performance in the final evaluation phase. This experience clearly highlights the critical importance of the proper model cross-validation however costly and complex it may be in the case of hybrid, ensemble and stacked models.

VIII. CONCLUSIONS

In this competition, we aimed to improve the predictive capabilities of the already effective models in the gradient boosting family, specifically XGBoost and LightGBM. To meet this challenge, we utilized various GBDT techniques with diverse ensemble strategies, achieving enhanced performance by aggregating a broader array of model variants. Further-

more, we implemented regression-based stacking and carefully selected the top-performing ensemble models, prioritizing a balance between performance and diversity to optimize results. This approach allowed us to refine and advance the effectiveness of our predictive models within the ensemble framework.

During the initial phase of the competition, our hybrid model showed promising results, securing a leading score of 0.5941, which placed us at the forefront among all competitors. However, this early success was not sustained in the final evaluation phase, where the model's performance fell to a score of 0.8415. This significant drop in performance underscores potential issues with the model's stability and its ability to generalize effectively across different testing environments, transitioning from a controlled setting to one that is more diverse and unpredictable. This experience highlights the inherent challenges in financial predictive modeling and paves the way for future research aimed at developing more robust and resilient AI-driven trading systems.

REFERENCES

- [1] S. Yadav and K. P. Sharma, "Statistical Analysis and Forecasting Models for Stock Market," *2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC)*, Jalandhar, India, 2018, pp. 117-121, doi: 10.1109/ICSCCC.2018.8703324.
- [2] J. Bagul, P. Warkhade, T. Gangwal and N. Mangaonkar, "ARIMA vs LSTM Algorithm – A Comparative Study Based on Stock Market Prediction," *2022 5th International Conference on Advances in Science and Technology (ICAST)*, Mumbai, India, 2022, pp. 49-53, doi: 10.1109/ICAST55766.2022.10039560.
- [3] R. Karim, M. K. Alam and M. R. Hossain, "Stock Market Analysis Using Linear Regression and Decision Tree Regression," *2021 1st International Conference on Emerging Smart Technologies and Applications (eSmarTA)*, Sana'a, Yemen, 2021, pp. 1-6, doi: 10.1109/eSmarTA52612.2021.9515762.
- [4] Z. Liu, Z. Dang and J. Yu, "Stock Price Prediction Model Based on RBF-SVM Algorithm," *2020 International Conference on Computer Engineering and Intelligent Control (ICCEIC)*, Chongqing, China, 2020, pp. 124-127, doi: 10.1109/ICCEIC51584.2020.00032.
- [5] Y. Wei and V. Chaudhary, "The Directionality Function Defect of Performance Evaluation Method in Regression Neural Network for Stock Price Prediction," *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, Sydney, NSW, Australia, 2020, pp. 769-770, doi: 10.1109/DSAA49011.2020.00108.
- [6] M. Faraz, H. Khaloozadeh and M. Abbasi, "Stock Market Prediction-by-Prediction Based on Autoencoder Long Short-Term Memory Networks," *2020 28th Iranian Conference on Electrical Engineering (ICEE)*, Tabriz, Iran, 2020, pp. 1-5, doi: 10.1109/ICEE50131.2020.9261055.
- [7] J. Creighton and F. H. Zulkernine, "Towards building a hybrid model for predicting stock indexes," *2017 IEEE International Conference on Big Data (Big Data)*, Boston, MA, USA, 2017, pp. 4128-4133, doi: 10.1109/BigData.2017.8258433.
- [8] A. Durgapal and V. Vimal, "Prediction of Stock Price Using Statistical and Ensemble learning Models: A Comparative Study," *2021 IEEE 8th Uttar Pradesh Section International Conference on Electrical, Electronics and Computer Engineering (UPCON)*, Dehradun, India, 2021, pp. 1-6, doi: 10.1109/UPCON52273.2021.9667644.
- [9] P. K. Aithal, U. D. Acharya, M. Geetha, R. Sagar and R. Abraham, "A Comparative Study of Deep Neural Network and Statistical Models for Stock Price Prediction," *2022 3rd International Conference for Emerging Technology (INCET)*, Belgaum, India, 2022, pp. 1-5, doi: 10.1109/INCET54531.2022.9824487.
- [10] X. Zheng, J. Cai and G. Zhang, "Stock Trend Prediction Based on ARIMA-LightGBM Hybrid Model," *2022 3rd Information Communication Technologies Conference (ICTC)*, Nanjing, China, 2022, pp. 227-231, doi: 10.1109/ICTC55111.2022.9778304.
- [11] R. Jaiswal and B. Singh, "A Hybrid Convolutional Recurrent (CNN-GRU) Model for Stock Price Prediction," *2022 IEEE 11th International Conference on Communication Systems and Network Technologies (CSNT)*, Indore, India, 2022, pp. 299-304.
- [12] Aleksandar M. Rakicevi, Pavle D. Milosevic, Ivana T. Dragovic, Ana M. Poledica, Milica M. Zukanovic, Andrzej Janusz, Dominik Slezak: Predicting Stock Trends Using Common Financial Indicators: FedCSIS 2024 Data Science Challenge on KnowledgePit.ai Platform, *Proceedings of the 19th Conference on Computer Science and Intelligent Systems (FedCSIS)*, 2024.
- [13] L. Mason, J. Baxter, P.L. Bartlett, and M. Frean. Boosting Algorithms as Gradient Descent In S.A. Solla and T.K. Leen and K. Müller. *Advances in Neural Inf. Processing Sys.* 12: 512–518, MIT Press, 1999.
- [14] J.H. Friedman. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* 29(5): 1189-1232, 2001.
- [15] M. Liu, L. Cen and D. Ruta, "Gradient Boosting Models for Cybersecurity Threat Detection with Aggregated Time Series Features," *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, Warsaw, Poland, 2023, pp. 1311-1315, doi: 10.15439/2023F4457.
- [16] D. Ruta, M. Liu and L. Cen, "Beating Gradient Boosting: Target-Guided Binning for Massively Scalable Classification in Real-Time," *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, Warsaw, Poland, 2023, pp. 1301-1306, doi: 10.15439/2023F7166.
- [17] D. Ruta, M. Liu, L. Cen. Feature Engineering for Predicting Frags in Tactical Games. *Proc. Int. Conf. 2023 IEEE International Conference on Multimedia and Expo*, 2023.
- [18] D. Ruta, M. Liu, L. Cen and Q. Hieu Vu. Diversified gradient boosting ensembles for prediction of the cost of forwarding contracts. *Proc. Int. 17th Conf. on Computer Science and Intelligence Systems*, 2022.
- [19] Q. Hieu Vu, L. Cen, D. Ruta and M. Liu. Key Factors to Consider when Predicting the Costs of Forwarding Contracts. *Proc. Int. Conf. 2022 17th Conf. on Computer Science and Intelligence Systems*, 2022.
- [20] D. Ruta, L. Cen, M. Liu and Q. Hieu Vu. Automated feature engineering for prediction of victories in online computer games. *Proc. Int. Conf on Big Data*, 2021.
- [21] Q. Hieu Vu, D. Ruta, L. Cen and M. Liu. A combination of general and specific models to predict victories in video games. *Proc. Int. Conf. on Big Data*, 2021.
- [22] D. Ruta, L. Cen and Q. Hieu Vu. Deep Bi-Directional LSTM Networks for Device Workload Forecasting. *Proc. 15th Int. Conf. Comp. Science and Inf. Sys.*, 2020.
- [23] L. Cen, D. Ruta and Q. Hieu Vu. Efficient Support Vector Regression with Reduced Training Data. *Proc. Fed. Conf. on Comp. Science and Inf. Sys.*, 2019.
- [24] D. Ruta, L. Cen and Q. Hieu Vu. Greedy Incremental Support Vector Regression. *Proc. Fed. Conf. on Computer Science and Inf. Sys.*, 2019.
- [25] Q. Hieu Vu, D. Ruta and L. Cen. Gradient boosting decision trees for cyber security threats detection based on network events logs. *Proc. IEEE Int. Conf. Big Data*, 2019.
- [26] L. Cen, A. Ruta, D. Ruta and Q. Hieu Vu. Regression networks for robust win-rates predictions of AI gaming bots. *Int. Symp. Advances in AI and Apps (AAIA)*, 2018.
- [27] Q. Hieu Vu, D. Ruta, A. Ruta and L. Cen. Predicting Win-rates of Hearthstone Decks: Models and Features that Won AAIA'2018 Data Mining Challenge. *Int. Symp. Advances in Artificial Intelligence and Apps (AAIA)*, 2018.
- [28] L. Cen, D. Ruta and A. Ruta. Using Recommendations for Trade Returns Prediction with Machine Learning. *Int. Symp. on Methodologies for Intelligent Sys. (ISMIS)*, 2017.
- [29] A. Ruta, D. Ruta and L. Cen. Algorithmic Daily Trading Based on Experts' Recommendations. *Int. Symp. on Methodologies for Intelligent Systems (ISMIS)*, 2017.
- [30] Q. Hieu Vu, D. Ruta and L. Cen. An ensemble model with hierarchical decomposition and aggregation for highly scalable and robust classification. *12th Int. Symp. Advances in AI and Applications (AAIA)*, 2017.
- [31] L. Cen and D. Ruta. A Map based Gender Prediction Model for Big E-Commerce Data. *The 3rd IEEE Int. Conf. on Smart Data*, 2017.
- [32] D. Ruta and L. Cen. Self-Organized Predictor of Methane Concentration Warnings in Coal Mines. *Proc. Int. Joint Conf. Rough Sets, LNCS*, Springer, 2015.
- [33] <https://machinelearningmastery.com/hyperparameter-optimization-with-random-search-and-grid-search/>.
- [34] <https://developers.google.com/machine-learning/crash-course/classification/roc-and-auc/>.

Forecasting Stock Trends with Feedforward Neural Networks

Marcin Traskowski
University of Warsaw
Warsaw, Poland

Email: traskowski.marcin@gmail.com

Eyad Kannout
University of Warsaw
Warsaw, Poland

Email: eyad.kannout@mimuw.edu.pl

Abstract—Stock market prediction stands as a complex and crucial task, pivotal for enhancing the overall stability and efficiency of financial markets by offering essential insights into market movements and trends. In this study, we introduce a simple yet potent model based on feedforward neural networks to tackle this challenge effectively. Our approach leverages advancements in machine learning and deep learning to analyze large datasets of financial statements, demonstrating promising results in forecasting stock trends.

Index Terms—Stock Market Prediction, Neural Networks, Deep Learning, Forecasting, Classification

I. INTRODUCTION

FORECASTING stock trends has long been an intriguing and challenging problem for researchers and enthusiasts in the fields of finance and data science. Accurate predictions can lead to significant financial gains and help investors make informed decisions. The factors and sources of information to consider are numerous and diverse, making it very challenging to foresee future stock market behavior accurately [1]. It is clear that precise prediction of stock prices is elusive. Fama et al. [2] introduced the efficient market hypothesis, which asserts that an asset's current price always reflects all available prior information. Additionally, the random walk hypothesis, introduced by Burton [3], proposes that a stock's price movements are independent of its past, implying that tomorrow's price will rely exclusively on tomorrow's information, irrespective of today's price. Together, these hypotheses suggest that accurately predicting stock prices is impossible. Nonetheless, extensive research has been conducted to address this issue, proposing a range of methodologies across multiple disciplines, including economics, statistics, physics, and computer science [4].

Traditional methods of stock trend analysis often rely on statistical models, such as ARIMA [5], which may not effectively capture the complex, non-linear, and often unstable patterns present in financial data. In recent years, advancements in machine learning and deep learning have opened new avenues for analyzing and predicting stock trends [6]. Neural networks, in particular, have shown promise due to their ability to learn and generalize from large datasets. This paper presents a simple yet effective neural network approach to predicting stock trends, developed as part of the FedCSIS 2024 Data Mining Challenge ¹ [7].

¹<https://knowledgepit.ai/fedcsis-2024-challenge/>

The paper is organized as follows: the next part reviews related work in stock predictions, followed by a description of the competition data and preprocessing methods. Next, we detail the model architecture, hyperparameters and evaluation metrics. Subsequently, we present the experimental results. The last part is reserved for final conclusions and observations.

II. RELATED WORKS

Recent advancements in machine learning and deep learning have significantly enhanced the ability to predict stock market trends. Numerous studies have explored different approaches and models, showcasing varying degrees of success and innovation.

Random Forest, an ensemble learning method, has been widely used for its robustness and accuracy in stock market prediction. By constructing multiple decision trees and aggregating their results, Random Forest reduces overfitting and improves generalization. Its ability to handle large datasets with high dimensionality makes it particularly effective for financial market predictions [8].

Other research endeavors have concentrated on employing support vector machines (SVMs) to improve the accuracy of stock market predictions through categorization of examples. SVM models represent examples as points in a multidimensional space, with the objective of maximizing the separation between different categories. New examples are then classified based on the category they are most likely to belong to [9]. Liu et al. [10] developed a model using the RBF-SVM algorithm to enhance stock price prediction, accurately assess short-term stock price movements, and provide more reliable guidance for stock market analysis and investor decision-making.

Another prominent method involves using Recurrent Neural Networks (RNNs) and Long Short-Term Memory networks (LSTMs) due to their ability to handle sequential data [11]. Research has shown that LSTMs can outperform traditional machine learning models in predicting stock price movements by capturing temporal dependencies in financial data [12].

The application of Convolutional Neural Networks (CNNs) to extract spatial features from stock market data has also shown promise. Research has utilized hybrid models combining CNNs and LSTMs to predict stock prices, showing improved performance over standalone models [13].

Attention mechanisms, particularly through Transformer models, have demonstrated significant potential in stock trend prediction. By effectively handling long-term dependencies and integrating both technical indicators and sentiment analysis, Transformer-based models like StockFormer [14] have shown superior predictive capabilities compared to traditional RNN-based models. Furthermore, stock market prices are significantly influenced by the sentiments of stakeholders, which can be assessed through the analysis of news, social media data, and other indicators. Kumar et al. [15] explored the correlation between the sentiment polarity of news and a company's stock price. They employed the SVM-LSTM-GRU Composite Model to predict a company's stock price based on news related to that company.

The research literature has documented numerous efforts to develop recommendation systems [16], [17] that advise users on whether to buy or sell stocks [18]. Additionally, Association Rule Mining (ARM) is widely used to create recommender systems [19]. In stock predictions, several systems exist for monitoring and predicting stock prices, but they typically focus on individual stocks and do not account for the inter-relationships between stocks or their connections with the stock market index. Paranjape-Voditel et al. [20] employed various ARM techniques—such as fuzzy ARM, weighted fuzzy ARM, ARM with time lags, fuzzy ARM with time lags, and weighted fuzzy ARM with time lags—to predict relationships between stocks. This approach forms the basis for portfolio management and provides recommendations for mutual funds.

These studies collectively indicate that leveraging advanced deep learning architectures, ensemble methods, and attention mechanisms can significantly enhance the predictive accuracy of stock market models.

III. DATASETS

The competition included training and test sets, consisting of 8,000 and 2,000 examples, respectively. Each example represents a financial statement announcement for one of the chosen 300 companies. Each record consists of 119 elements. The first column is a categorical value that represents the company's sector. The next 58 columns contain values for key financial indicators. The following 58 columns represent the 1-year (absolute) change for each of these previous financial indicators. The last two columns are targets:

- 1) The target column "Class" can have three possible values:
 - -1: "sell" (do not invest)
 - 0: "hold"
 - 1: "buy" (invest)
- 2) The target column "Perform" gives the value of risk-return performance for a period after the announcement. Those values range between -0.5 and 0.5. Small negative values correspond to the "sell" class, large positive values correspond to the "buy" class, and values close to 0 correspond to the "hold" class.

IV. DATA PREPROCESSING

Before starting the training, we needed to address a few existing issues with the datasets.

A. Missing Values

There are two distinct types of missing values in the provided datasets:

- 1) Non-available: To fill these empty spaces with values that minimize the difference from actual unknown data, we opted to use the mean of the column values.
- 2) Non-applicable: We set these values to 0. Using any other value might suggest to our neural network model that these components should be active.

These two types were marked differently. Non-available data were represented as empty strings, while non-applicable data were marked with the "NA" string. This distinction allowed us to easily differentiate between the two cases.

B. Categorical Values

As mentioned earlier, the first column is a categorical variable indicating the company's sector, with eleven possible values. To correctly process this data, we one-hot encoded this column. This increased the number of columns in our datasets by ten, resulting in a total of 129 columns.

C. Standardization

To improve the performance of our model, we decided to normalize the datasets using z-score normalization [21], which calculates the standard score for each feature. The standard score, or z-score, of a feature x is calculated as:

$$z = \frac{x - \mu}{\sigma}$$

where μ is the mean of the feature values in the training set, and σ is the standard deviation of the feature values in the training set.

By applying z-score normalization, each feature will have a mean of 0 and a standard deviation of 1 in the training set. This transformation ensures that all features are on a similar scale, which can help improve the convergence speed of optimization algorithms and prevent features with larger scales from dominating the learning process.

The scaling parameters (mean and standard deviation) computed on the training set are then used to transform the test set, ensuring consistency in scaling across both datasets.

V. MODEL ARCHITECTURE

We will test three neural networks with varying depths to see how well each of them manages the prediction task. All models will share common hyperparameters:

- Activation function: Hyperbolic Tangent (Tanh)
- Number of epochs: 100
- Learning rate: 0.00005
- Batch size: 40
- Optimizer: AdamW
- Loss function: mean squared error (MSE)




Fig. 1. Histogram of values from column "Perform" from Test Dataset

- Dropout probability: 20%
- Hidden layers size: 256

Each model accepts an input vector of 127 dimensions, which is processed through multiple dense hidden layers with dropout for regularization. To improve performance, we employed Xavier Glorot initialization [22] for our models weights. It works by initializing the weights using the formula:

$$W \sim \mathcal{U} \left(-\sqrt{\frac{6}{n_{in} + n_{out}}}, \sqrt{\frac{6}{n_{in} + n_{out}}} \right)$$

where n_{in} and n_{out} are the size of input and output units in the layer.

Batch normalization [23] is applied after each layer to stabilize training. As mentioned earlier, the models vary in terms of the number of hidden layers. Specifically, the first model consists of 2 hidden layers, the second model includes 6 hidden layers and the third model incorporates 10 hidden layers. These configurations were chosen to explore how increasing depth impacts the model’s ability to learn and generalize from the data.

The final layer of each model predicts a single value corresponding to the risk-return performance (target column "Perform"). Given that these performance values typically range between -0.5 and 0.5, we opted for the Hyperbolic Tangent activation function because it confines outputs to the range [-1, 1]. For a clearer visualization of these values, Fig 1 illustrates their distribution.

Our models are well-suited for the regression task of predicting risk-return performance. To adapt them for a classification problem, where each input should be assigned to one of the three classes described in previous sections, we implement the following post-processing steps:

We identify the highest value of risk-return performance from train dataset that still belongs to the class -1, marking this value as low_limit . Similarly, we identify the smallest value of risk-return performance from train dataset that corresponds to the class 1, marking this value as $upper_limit$. The classification proceeds as follows:

- if model predicts value that is less than low_limit , then we classify this instance as class -1
- if it predicts value from range $[low_limit, upper_limit]$, then we classify it as class 0

- in case we get a value bigger than $upper_limit$, then we classify it as class 1

The values of these boundaries rounded to five decimal places are:

$$low_limit = -0.01504$$

$$upper_limit = 0.04008$$

VI. EXPERIMENTAL RESULTS

In this section, we present the results of our experiments. The models are listed in the order they were introduced in the previous sections (Model 1 has 2 hidden layers, Model 2 has 6 hidden layers and Model 3 has 10 hidden layers).

A. Evaluation metrics

Our model can be assessed using two metrics. The first metric is the mean absolute error (MAE) [24] between the predicted values of our model and the true values of risk-return performance. We opt for MAE over mean squared error (MSE) due to the small magnitude of values involved in our predictions. While MSE could provide useful insights as well, we prioritize MAE for its straightforward interpretation, especially when presenting results with fewer decimal points. The MAE is calculated as the average of the absolute differences between the true values y_i and the predicted values \hat{y}_i :

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

The second metric, which we will call classification weighted error, corresponds to the evaluation criterion used in the FedCSIS 2024 Data Mining Challenge. After predicting classes, a confusion matrix is constructed. The evaluation involves calculating the average error cost using a predefined error cost matrix ²:

	-1	0	1
-1	0	1	2
0	1	0	1
1	2	1	0

The classification weighted error is determined using the following formula:

$$\frac{1}{\text{length}} \sum (\text{cost matrix} \circ \text{confusion matrix})$$

Here, \circ denotes the Hadamard product (element-wise multiplication), the term "length" represents the total number of predictions (2000 in our case) and the sum is over all the elements of the resulting matrix.

This method imposes a higher penalty for misclassifications between classes 1 and -1, specifically when predicting 'invest' for a true class of 'do not invest', and vice versa.

B. Achieved Scores

The table I displays the scores achieved for the two aforementioned metrics.

²<https://knowledgepit.ai/fedcsis-2024-challenge/>




Fig. 2. Histograms of differences between true values of risk-return performance and predicted values




Fig. 3. Confusion matrices for all three models

TABLE I
ACHIEVED SCORES ROUNDED TO FOURTH DECIMAL POINT

Model	MAE	Classification Weighted Error
Model 1	0.1174	0.861
Model 2	0.1158	0.839
Model 3	0.1197	0.885

Fig 2 presents histograms of differences between true values and predicted values. Although the distributions appear quite similar, these slight variances lead to marginally different outcomes in those two aforementioned metrics.

In Fig 3, the confusion matrices for each model are presented. We observe consistent patterns across all three models, such as managing to correctly classify a large number of '1's while also frequently misclassifying classes -1 and 1 as 0. Although adjusting the values of previously mentioned *low_limit* and *upper_limit* might seem like a logical step to reduce the frequency of classifying instances as 0, this adjustment, most of the times, does not significantly improve the weighted classification error. This is because the error metric penalizes misclassifications between classes -1 and 1 more severely. Therefore, by allowing the model to maintain such a buffer of class "0", we are empirically able to achieve

better results in terms of weighted classification error.

Although our models achieved comparable results, it is evident that hyperparameter selection is a crucial aspect of our method. Analyzing the outcomes, we observed that the smallest model slightly underfitted while the largest one overfitted. Therefore, despite each model achieving a fairly good level of performance, fine-tuning them requires a thorough grid search.

VII. CONCLUSIONS

In this study, we proposed a robust framework for predicting stock market trends using feedforward neural networks. We aimed to enhance the accuracy of forecasting models by leveraging advancements in deep learning and extensive preprocessing techniques. Our approach involved tackling challenges such as handling large datasets of financial statements, managing missing data, and preprocessing categorical variables and numerical features. We explored three neural network architectures with varying numbers of hidden layers, evaluating their performance based on both regression (mean absolute error) and classification (weighted error) metrics. This approach achieved best cumulative performance in terms of the risk-return aspect in FedCSIS 2024 Data Mining Challenge.

REFERENCES

- [1] P. Tran, P. Anh, P. Tam, and C. Nguyen, "Applying machine learning algorithms to predict the stock price trend in the stock market – the case

- of vietnam,” *Humanities and Social Sciences Communications*, vol. 11, 03 2024.
- [2] E. F. Fama, “Efficient capital markets: A review of theory and empirical work,” *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970.
- [3] N. Burton, *An Analysis of Burton G. Malkiel’s A Random Walk Down Wall Street*. Macat Library, 01 2018.
- [4] M. Agrawal, A. Khan, and P. Shukla, “Stock price prediction using technical indicators: A predictive model using optimal deep learning,” *International Journal of Recent Technology and Engineering*, vol. 8, pp. 2297–2305, 07 2019.
- [5] W. R. Kinney, “Arima and regression in analytical review: An empirical test,” *The Accounting Review*, vol. 53, no. 1, pp. 48–60, 1978.
- [6] G. Sonkavde, D. S. Dharrao, A. M. Bongale, S. T. Deokate, D. Doreswamy, and S. K. Bhat, “Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications,” *International Journal of Financial Studies*, vol. 11, no. 3, 2023.
- [7] A. M. Rakicevic, P. D. Milosevic, I. T. Dragovic, A. M. Poledica, M. M. Zukanovic, A. Janusz, and D. Slezak, “Predicting stock trends using common financial indicators: A summary of fedcsis 2024 data science challenge held on knowledgepit.ai platform,” in *Proceedings of FedCSIS 2024*, 2024.
- [8] J. Zheng, D. Xin, Q. Cheng, M. Tian, and L. Yang, “The random forest model for analyzing and forecasting the us stock market in the context of smart finance,” 2024.
- [9] T. Strader, J. Rozycki, T. Root, and Y.-H. Huang, “Machine learning stock market prediction studies: Review and research directions,” *Journal of International Technology and Information Management*, vol. 28, pp. 63–83, 01 2020.
- [10] Z. Liu, Z. Dang, and J. Yu, “Stock price prediction model based on rbf-svm algorithm,” in *2020 International Conference on Computer Engineering and Intelligent Control (ICCEIC)*, pp. 124–127, 2020.
- [11] S. Mehtab, J. Sen, and A. Dutta, “Stock price prediction using machine learning and lstm-based deep learning models,” in *Machine Learning and Metaheuristics Algorithms, and Applications* (S. M. Thampi, S. Piramuthu, K.-C. Li, S. Berretti, M. Wozniak, and D. Singh, eds.), (Singapore), pp. 88–106, Springer Singapore, 2021.
- [12] T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” *European Journal of Operational Research*, vol. 270, no. 2, pp. 654–669, 2018.
- [13] J. Eapen, D. Bein, and A. Verma, “Novel deep learning model with cnn and bi-directional lstm for improved stock market index prediction,” in *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, pp. 0264–0270, 2019.
- [14] H. Kaeley, Y. Qiao, and N. Bagherzadeh, “Support for stock trend prediction using transformers and sentiment analysis,” 2023.
- [15] R. Kumar, C. M. Sharma, V. M. Chariar, S. Hooda, and R. Beri, “Emotion analysis of news and social media text for stock price prediction using svm-lstm-gru composite model,” in *2022 International Conference on Computational Intelligence and Sustainable Engineering Solutions (CISES)*, pp. 329–333, 2022.
- [16] E. Kannout, M. Grzegorowski, and H. Son Nguyen, *Toward Recommender Systems Scalability and Efficacy*, pp. 91–121. Cham: Springer International Publishing, 2023.
- [17] E. Kannout, M. Grzegorowski, M. Grodzki, and H. S. Nguyen, “Clustering-based frequent pattern mining framework for solving cold-start problem in recommender systems,” *IEEE Access*, vol. 12, pp. 13678–13698, 2024.
- [18] Sharma, Vikram, Rakhra, Manik, and Mathur, Gauri, “Hybrid approaches for stocks prediction and recommendation system,” *E3S Web Conf.*, vol. 453, p. 01047, 2023.
- [19] E. Kannout, H. S. Nguyen, and M. Grzegorowski, “Speeding up recommender systems using association rules,” in *Intelligent Information and Database Systems* (N. T. Nguyen, T. K. Tran, U. Tukayev, T.-P. Hong, B. Trawiński, and E. Szczerbicki, eds.), (Cham), pp. 167–179. Springer Nature Switzerland, 2022.
- [20] P. Paranjape-Voditel and U. Deshpande, “An association rule mining based stock market recommender system,” in *2011 Second International Conference on Emerging Applications of Information Technology*, pp. 21–24, 2011.
- [21] N. Fei, Y. Gao, Z. Lu, and T. Xiang, “Z-score normalization, hubness, and few-shot learning,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 142–151, 2021.
- [22] Y. Bengio and X. Glorot, “Understanding the difficulty of training deep feed forward neural networks,” *International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 01 2010.
- [23] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” 2015.
- [24] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, vol. 2. Springer, 2009.

Author Index

- Abbasi, Ali 279
Abbas, Musarat 555, 615
Ajwani, Deepak 689
Akalin, Özgün 207
Albrecht, Jens 501
Alexandre, Rosana 655
Alghamdi, Ahmad 63
Alonso, Marco 301
Alptekin, Gülfem Isiklar 207
Alqahtani, Mohammed 75
Alshahrani, Fayez 87
Alshammari, Wafa 213
Alves, Filipe 279
Alves, Manuel 655
Alzahrani, Amal 531
Amonov, Bekzod 739
Andresel, Medina 219
- B**
Babuc, Diogen 537
Bacco, Manlio 543
Bădică, Costin 225
Bajrami, Merxhan 107
Barone, Marco 549, 573
Barsocchi, Paolo 543
Batool, Farwa 555, 615
Bavčar, Urban 375
Beloff, Natalia 63, 75, 87, 213, 343, 531
Bělohoubek, Marek 567
Bicevska, Zane 289
Bicevskis, Janis 289
Bieniasz, Jędrzej 621
Bierska, Adela 195
Bjeladinović, Srđa 395
Bolwerk, Twan 301
Bondaruk, Łukasz 579
Boschetti, Mirco 475
Bratskas, Romaio 561
Brunori, Gianluca 543
Buck, Luisa 751
Buhnova, Barbora 119, 195
Bulcão-Neto, Renato F. 187
Burger, Florian 385
Butler, Jethro 513
Buyya, Rajkumar 21
Bylina, Jarosław 661
- Č**
Čabarkapa, Radoš 375
Candiani, Gabriele 475
Carvalho, Pedro 279
- Cen, Ling 761
Challenger, Moharram 637
Charvát, Karel 485
Chessa, Stefano 543
Ciaschi, Matteo 549, 573
Cisotto, Giulia 313
Colim, Ana 655
Costa, Joao Pita 375
Crivello, Antonino 543
Czyżnikiewicz, Mateusz 579
- D**
Dahrouje, George 561
Damaševičius, Robertas 99, 141
Degórski, Łukasz 579
Díaz-Salazar, Aldo A. 187
Diebelis, Edgars 289
Dimitri, Giovanna Maria 543
Dobрева, Jovana 107
Doerpinghaus, Jens 417
Dora, Shirin 701
Dragović, Ivana T. 731
Dujmović, Jozo 1
- E**
Egger, Manuel 53
Engels, Stefan 585
- F**
Faia, Ricardo 279
Faria, Pedro 279
Faruga, Michael Jan 131
Fauster, Norman FRM. 53
Fazzioni, Daniel 187
Fidalgo, Pablo de Juan 591
Fidanova, Stefka 597
Filipovska, Elena 107
Fiorino, Mario 615
Fischer, Alexander 319
Franczyk, Bogdan 157
Frank, Ulrich 9
Freitas, Tiago Carvalho. 331
Fung, Yat Chun 739
Furier, Marc 751
Furtula, Filip 395
- G**
Gallo, Ignazio 475
Ganzha, Maria 597
Gerber, Philippe 513
Ghiwaa, Taif 343
Göktepe, Okan Mert 751
Gori, Marco 543

Goupil, Alban	707	Kocian, Alexander	543
Gradoń, Kacper	603	Kolagari, Ramin Tavakoli	319
Gross, Ty	745	Koller, Linda	131
Grużewski, Mateusz	671	Kosar, Tomáš	395
Gryka, Paweł	603	Kovalenko, Alexander	485
Guimarães, Murilo O.	187	Kozak, Mateusz	237
Gul, Seren	631	Kozłowski, Marek	603
H abarta, Filip	725	Křen, Jan	683
Hajdu, László	609	Kubiak, Jakub	579
Hamlaoui, Mahmoud El	249, 443	Kubis, Marek	579
Hasan, Tehreem	555, 615	Kuchařková, Kateřina	683
Hein, Kristine	355	Kutyła, Miłosz	603
Heller, Thomas	157	Kvapil, Jiří	485
Henriques, Pedro Rangel	331, 453	L ameski, Petre	107
Herrera, Francisco	45	Lam, Phat	219
Hillman, Vellislava	107	Lange, Berit	53
Holzheuser, Anna	253	Laskov, Lasko	405
Horvadvoska, Olga	53	Lautenbach, Max	751
Hościłowicz, Jakub	621	Leber, Žiga	395
Humm, Bernhard G.	253	Leventakis, George	561
Hutzler, Guillaume	363	Leyh, Christian	131, 231
I gnaciuk, Przemysław	265	Li, Baihua	701
Ignesti, Giacomo	625	Liberti, Leo	677
Ilager, Shashikant	21	Lin, Chang	755
Iqbal, Danish	119	Liu, Ming	761
Ivanov, Veselin	597	Ljubisavljević, Miloš	395
J anicki, Artur	603, 621	London, András	609
Janusz, Andrzej	731	Lorenz, Alisa	131
Jejić, Olga	395	Lotti, Edoardo	313
Jenko, Jakob	375	Lübbert, Daniel	385
Jovanovic, Jelena	35	Łuczak, Łukasz Piotr	265
Judák, Jusztna	751	Lukas, Vojtěch	683
Junior, Iwens G. Sene	187	Luković, Ivan	395
K alan, Reza Shokri	631	Luţan, Elena-Ruxandra	225
Kamilaris, Andreas	643	M akajić-Nikolić, Dragana	149
Kammeyer, Alexander	385	Malá, Ivana	725
Kannout, Eyad	767	Malliavin, Therese	677
Kaplánek, Tomáš	683	Mangroliya, Meetkumar Pravinbhai	417
Karaduman, Burak	637	Manzoni, Sara L.	313
Karatsiolis, Savvas	643	Marek, Luboš	725
Kardas, Geylani	169, 637	Marinov, Marin	405
Kawulok, Michal	259	Martinelli, Massimo	625
Kesarovski, Todor	713	Maskeliunas, Rytis	99
Kesselbacher-Pirker, Max	53	Maskeliūnas, Rytis	141
Khan, Imran	343	Mathie, Janette	701
Kimmich, Christian	53	Mernik, Marjan	395
Kirilov, Leoneed	597	Mihajlov, Teodora	649
Kitanović, Olivera	465, 649	Milić, Miloš	149
Klaudiel, Hanna	363	Miller, Gloria J.	429
Klaudiel, Witold	363	Milosević, Pavle D.	731
Klein, Olivier	513	Mladenovska, Ana	107
		Moreira, Filipe	655
		Moroni, Davie	625

Moshkov, Mikhail	667	Rodiah, Isti	53
Mothes, Carina	501	Rodrigues, Nelson	655
Motii, Anas	443	Rodrigues, Ricardo	279
Mucherino, Antonio	677	Roß, Julian	745
Müller, Julia	231	Rudkowski, Jan	621
Münker, Gregor	751	Rudolph, Eric	501
N assar, Mahmoud	249	Ruta, Dymitr	761
Nešić, Milica Ikonić	465, 649	S aritas, Hidayet Burak	169
Neto, Alvaro Costa	331, 453	Savvidis, Petros	561
Netscher, Maike	493	Schäffer, Thomas	231
Nguyen, Tin	219	Schauer, Stefan	53
Nicolaou, Nicolas	643	Schindler, Alexander	219
Niedziółka-Domański, Rafał	661	Schuts, Mathijs	301
Nikolić, Nebojša	243	Sconfienza, Umberto	513
O liveira, João	655	Seer, Hanna	501
Ostojčić, Dragana	395	Seidl, Michael	219
Ostonov, Azimkhon	667	Sengupta, Aditya	525
P alkowski, Marek	671	Silva, Rosa Mariana	655
Panagi, Pieris	643	Šimánek, Petr	485
Pańkowska, Małgorzata	237	Škembarević, Milica	395
Panon, Marie-Laure	707	Skórzewski, Paweł	579
Papachristos, Dimitrios	561	Ślęzak, Dominik	731
Parigi, Lorenzo	475	Slezák, Vojtěch	683
Pasic, Aljosa	513, 591	Spinczyk, Dominik	237
Pereira, João	655	Stanković, Ranka	465, 649
Pereira, Maria João Varanda	331, 453	Stanojević, Bogdana	243
Pergl, Robert	567	Stanojević, Milan	243
Perrin, Eric	707	Stavrakis, Efsthios	643
Petalinkar, Saša	465	Štěpánek, Lubomír	719, 725
Petiz, Maria	279	T amenaoul, Hamza	249
Pham, Lam	219	Tang, Hieu	219
Piccardi, Armando	549	Tayebi, Dena	689
Pihrt, Jiří	485	Tegegn, Dagmawi Delelegn	313
Pluhár, András	609	Tezel, Baris Tekin	637
Poledica, Ana M.	731	Tiplea, Ferucio Laurentiu	177
Pommereau, Franck	363	Tolvanen, Juha-Pekka	61, 319
Popiołek, Paweł	621	Tomaszek, Krzysztof	237
Poray, Julia	157	Toscano, Piero	475
Poudel, Apeksha	745	Traskowski, Marcin	767
Purina, Liva	289	Tretola, Giancarlo	555, 615
Q erama, Enea	561	U llah, Saeed	695
Qurthobi, Ahmad	99	Ullah, Zaib	549, 555, 695
R akićević, Aleksandar M.	731	Utvić, Miloš	465
Rataj, Artur	363	V ale, Zita	279
Ray, Saurabh	689	Vassiliades, Vassilis	643
Reguzzoni, Ivan	313	Vitória, Arthur Ricardo Sousa	187
Ribeiro, Rui	279	Vladušič, Daniel	375, 513
Riedinger, Constanze	493	Vojnar, Daniel	195
Rocha, Wagner Da	677	Vrabie, Valeriu	707
Rockenfeller, Robert	417	Vu, Quang Hieu	761

W aqas, Abdullah	695	Zara, Ana Laura S. A.	187
Werlang, Arthur Allebrandt	745	Zarzosa, Susana González	591
White, Martin	63, 75, 87, 213, 343, 531	Zdravevski, Eftim	107
Whittaker, Michael	701	Zender, Alexander	253
Wiącek, Adam	579	Zenz, Hannes	53
Wille, Robert	585	Zhang, Shurong	707
Wolter, Katinka	385	Zherdev, Nikolay	513
Würth, Stefanie	231	Zhivkov, Petar	713
X ing, Ruiming	701	Zimmermann, Stephan	493
Z ancanaro, Alberto	313	Zoppis, Italo F.	313
		Zukanović, Milica M.	731
		Zyrek, Maciej	259
		Żytniewski, Mariusz	237