

HYBRINFOX at CheckThat! 2024 - Task 1: Enhancing Language Models with Structured Information for Check-Worthiness Estimation

Notebook for the HYBRINFOX Team at CheckThat! 2024 - Task 1

Géraud Faye^{1,2,*}, Morgane Casanova³, Benjamin Icard⁴, Julien Chanson⁶, Guillaume Gadek¹, Guillaume Gravier³ and Paul Égré⁵

¹Airbus Defence and Space, France

²Université Paris-Saclay, CentraleSupélec, MICS, France

³Université de Rennes, CNRS, Inria, IRISA, France

⁴LIP6, Sorbonne Université, CNRS, France

⁵Institut Jean-Nicod, CNRS, ENS-PSL, EHESS, France

⁶Mondeca, France

Abstract

This paper summarizes the experiments and results of the HYBRINFOX team for the CheckThat! 2024 - Task 1 competition. We propose an approach enriching Language Models such as RoBERTa with embeddings produced by triples (*subject ; predicate ; object*) extracted from the text sentences. Our analysis of the developmental data shows that this method improves the performance of Language Models alone. On the evaluation data, its best performance was in English, where it achieved an F1 score of 71.1 and ranked 12th out of 27 candidates. On the other languages (Dutch and Arabic), it obtained more mixed results. Future research tracks are identified toward adapting this processing pipeline to more recent Large Language Models.

Keywords

Hybrid AI, Text Classification, Check-Worthiness, Fact-Checking, Language Models

1. Introduction

The recent democratisation of social media has given the users unprecedented access to information, with the possibility to contribute knowledge as well as to share personal views and opinions. By the same token, however, it has also offered misinformation new paths to propagate, sometimes with massive impact. Because of that, automated misinformation detection and automated fact-checking have become tasks of central interest in the data science community.

This paper deals with a specific aspect of fact-checking, namely with the problem of check-worthiness estimation, presented as Task 1 of the the broader CheckThat! 2024 workshop [1], concerned with information quality evaluation. Various works dealing with fact-checking operate under the assumption that the entirety of the claims, sentences, or articles in a dataset are *checkworthy* [2, 3, 4]. But this approach can be inefficient, and a useful preliminary step is the identification of which claims are check-worthy, and which are not. The notion of check-worthiness is complex. Some claims in a text are not check-worthy simply because they are not declarative sentences and do not report even potential facts (viz. questions). Others are not check-worthy because, while making declarative assertions, they make claims of no consequence. Conversely, a declarative sentence with potentially harmful consequences is one that ranks high on check-worthiness. Others, finally, may not be check-worthy when they simply

CLEF 2024: Conference and Labs of the Evaluation Forum, September 09–12, 2024, Grenoble, France

*Corresponding author.

✉ geraud.faye@centralesupelec.fr (G. Faye); morgane.casanova@irisa.fr (M. Casanova); benjamin.icard@lip6.fr (B. Icard); julien.chanson@mondeca.com (J. Chanson); guillaume.gadek@airbus.com (G. Gadek); guillaume.gravier@irisa.fr (G. Gravier); paul.egre@ens.psl.eu (P. Égré)

ORCID 0000-0002-2985-5964 (G. Faye); 0009-0005-4530-5646 (B. Icard); 0000-0002-9114-7686 (P. Égré)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

report subjective views that are not susceptible of verification proper. It is a non-trivial challenge, therefore, to determine which claims in a document are specifically check-worthy.

Check-worthiness is a recent task [5], mostly covered using Language Models [6]. In this paper, we propose an approach designed to leverage structured information from the text, in order to enhance the representation obtained with a Language Model. Because the task of check-worthiness estimation is related to fact-checking, it seems appropriate to identify facts from the text to help the model predict check-worthiness. By using both structured facts and Language Models embeddings, we obtained better results than when using Language Models alone, ranking 12th among 27 competing teams for English (with an F1 score of 71.1). Results were more mixed for the non-English languages represented in the test set: in Dutch our method ranked 8th out of 16 candidates (F1 score of 58.9), and in Arabic it ranked 10th out of 14 (F1 score of 51.9).

In Section 2, we open with a quick review of the state of the art on the task. In Section 3, we spell out the functioning of the proposed processing pipeline. Then, Section 4 discusses preliminary results obtained with the initial training data and presents the final submitted results. Finally, some elements on the evolution and future use of the proposed hybrid system are presented in Section 5.

2. Related work

As explained in the previous section, check-worthiness is a fairly recent task, first mentioned in 2015 [5]. Several datasets have been constructed, like the ClaimBuster dataset [7], or the datasets proposed at the CheckThat workshops since 2018 [8]. These datasets focus on two main types of contexts for the task:

- Classifying sentences from a political debate. These could be used to ease fact-checking during television political debates, on datasets such as [9].
- Classifying tweets. Because they are easily and widely shared online, check-worthiness is an important task for online discussions to avoid information manipulation.

Both of these categories are important, and a commonality between them is their short format. However, the scope of this task of check-worthiness can be widened so as to also include online press, with an eye to so-called “pink slime” news [10], encompassing longer texts whose truthfulness is questionable.

Among pioneering approaches to the task, we find methods such as ClaimRank [11], using traditional NLP methods (e.g. lemmatization, TF-IDF) to identify check-worthy claims.

More recent approaches take advantage of the Transformer layer and of pretrained Language Models such as BERT [12], RoBERTa [13] or XLNet [14]. The fortune of these approaches can be seen in the 2023 CheckThat! Task 1 overview paper [15], showing that nearly all teams used a transformer-based Language Model.

With the even more recent development of Large Language Models and of Generative AI, a natural shift has been made toward the use of LLMs, relying on prompt engineering [16] and in-context learning [17] to achieve check-worthiness estimation. These approaches were used by the winners of this year’s competition, both in English [18] and in Dutch [19].

3. Methodology

3.1. Model

A straightforward approach for check-worthiness estimation would be to use a pretrained Language Model fine-tuned with the provided training data. However, these language models produce embeddings that are opaque, even if they are sufficient most of the time. To increase the quality of the language model predictions, we propose to use them in conjunction with a small neural network able to leverage structured information from the input text. A visual description of the architecture is given in Figure 1. The processing pipeline is the following:

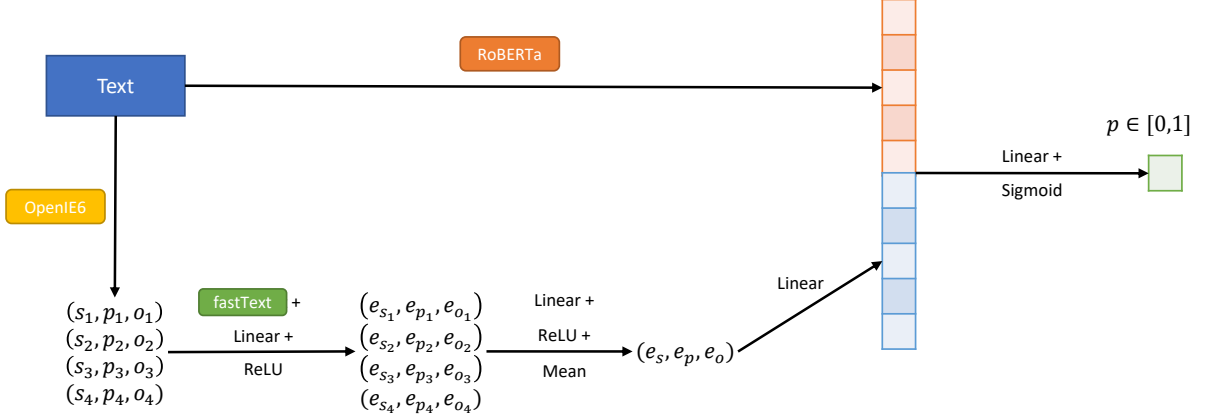


Figure 1: Our proposed architecture: adding structured information extracted from the text to enhance the LM embeddings. RoBERTa and OpenIE6 can be switched with other models for non-English languages.

1. To begin with, the text is embedded using a Language Model. In our case, the RoBERTa model [13] is used, producing an embedding of dimension 768. RoBERTa was chosen for its ease of use, its high performance for classification tasks and its relatively small size when compared to recent LLMs.
2. In parallel, the text is structured using an Open Information Extraction system. These systems extract information from the text in the form of triples (*subject; predicate; object*). We used OpenIE6 [20] to extract triples in the English language. Using triples allows us to produce structured information from the text and to reduce syntactic complexity, with the aim of helping sentence classification. A maximum limit of 4 triples by text were extracted, which is enough to consider all information triples for more than 90% of sentences. Each part of the triples is encoded using fastText [21], producing 3 vectors of dimension 300 per triple. These vector representations go through a dense layer with ReLU activation function. Then, they are averaged before being combined in the last layer to produce an embedding of dimension 768 (the same dimension as the Language Model).
3. Encodings from the previous two parts are concatenated and go through a dense layer with ReLU activation function with a final output producing the probability of being checkworthy by means of a sigmoid activation function.

The described architecture can be transposed to other languages when an OpenIE system and an LM are available. In the context of Task 1 of the evaluation, for Spanish, Dutch and Arabic, RoBERTa was swapped with a multilingual BERT [12].¹ The OpenIE6 system was replaced with Multi²OIE [22] in a zero-shot setting for non-English languages, providing worse performance than OpenIE6 on English, but allowing us to test the architecture on other languages. In principle, the same architecture can be used for any language, and in practice it is applicable to the 98 languages currently supported by the multilingual version of BERT.

3.2. Example

To better understand how this architecture works, we illustrate the pipeline with a simple example. We took a sentence from the training English dataset: *"I must remind him the Democrats have controlled the Congress for the last twenty-two years and they wrote all the tax bills."* This sentence comes from a debate between US presidential candidates Jimmy Carter and Gerald Ford on September 26th 1976. It is deemed to be checkworthy, as it contains allegations on Jimmy Carter's party.

¹<https://huggingface.co/google-bert/bert-base-multilingual-cased>

The classical NLP classification pipeline using RoBERTa would consist in producing a 768 dimension embeddings and doing classification passing this embeddings through a neural network. In our approach, we produce in parallel embeddings of the triples that can be extracted from the text. For this example, OpenIE6 produces the three following triples:

(I; must remind; him the Democrats have controlled the Congress for the last twenty-two years)

(the Democrats; have controlled; the Congress for the last twenty-two years)

(they; wrote; all the tax bills)

Only the last two triples contain the information that is check-worthy. However, there are no triple-level annotations, so we keep all triples. Firstly, we encode each subject, predicate and object with fastText, creating in this case three embedding vectors for each triple. These embeddings go through the same linear layer and the embeddings of subjects, predicates and objects are averaged by component (subject, predicate, object). This leads to three vectors of dimension 300, representing the subjects, predicates and objects of the triples extracted from the text. Finally they are concatenated and projected into a vector of dimension 768, the same dimension as the RoBERTa embeddings. RoBERTa embeddings and embeddings for triples are eventually concatenated for standard classification.

4. Results

This section is divided in three parts. The first presents the protocol of our model. The second part reports the evaluation of our proposed approach and of a standard Language Model, in order to measure how the additional triple processing part impacts performance. The third part contains an analysis of our submitted results, as well as of the difficulties encountered.

4.1. Training procedure

Each model was trained over 5 epochs on the train set. After each epoch, the model was evaluated on the dev set and the best model in terms of macro-F1 score was kept. The scores reported in Section 4.2 are the macro-F1 score on the dev-test set.

In our procedures, only the train set was used for training the model, the dev and dev-test sets being used for model selection. Reported results were produced by the models with the best dev-test macro-F1 score, which were also used to make predictions on the final test set.

4.2. Preliminary results on the development data

The main goal of our approach was to evaluate how combining structured information from the text with a standard Language Model is impacting performance. Results observed on the dev-test set provided before the competition are provided in Table 1.

Table 1

Macro-F1 scores on the dev-test split with a Language Model (RoBERTa⁽¹⁾ or bert-base-multilingual-cased⁽²⁾, where the superscripts indicate which LM was used for which language) and our own approach (LM+Triples). The F1-scores are multiplied by 100 to homogenize with the organizers' way of communicating scores.

	English ⁽¹⁾	Arabic ⁽²⁾	Dutch ⁽²⁾	Spanish ⁽²⁾
LM	84.042	58.273	40.866	59.975
LM+Triples	86.458	62.300	39.832	62.371
Performance gain	+2.416	+4.027	-1.034	+2.396

In most cases, our approach outperformed the LM baseline, achieving the highest performance gain in Arabic, followed by English and Spanish.

It appears that performance is generally lower for non-English languages. This is due to the fact that multilingual models perform worse, but they allowed us to process different languages with the same architecture and weights.

The same goes for the OpenIE system, with OpenIE6 being specifically trained on English, and Multi²OIE being used in a zero-shot setting (no non-English training sample was used for training). This limitation is pointed out in the Multi²OIE paper [22], but it is the only existing open-source OpenIE system able to process Arabic, Dutch, English and Spanish. As it relies on a multilingual BERT, it also suffers from lower performance for relatively low-resource languages, explaining the decrease in performance for Dutch.

4.3. Results on the evaluation data

The competition scores and ranking are shown in Table 2, with the scores of the best performing team (state-of-the-art for this dataset) and the baseline being also reported.

Table 2

Final results on the evaluation set. The reported F1-scores and ranking were provided by the organizers of the task.

Language	English	Arabic	Dutch
Best performance	80.2 (FactFinders [18])	56.9 (visty)	73.2 (TurQUaz [19])
Hybrinfox	71.1 (12/27)	51.9 (10/14)	58.9 (8/16)
Baseline	30.7 (27/27)	41.8 (13/14)	43.8 (14/16)

For the three languages, our approach outperformed the baseline by a substantial margin. Performance for non-English languages was mixed. The Arabic dataset proved to be challenging for all teams, with most candidate approaches getting scores between 50 and 55.

4.4. Discussion

While the proposed approach outperformed RoBERTa on the dev-test set, several upgrades could have been made to reduce possible errors in the processing pipeline. Firstly, it is well known that triples extracted with OpenIE are noisy and may not always contain useful facts for the task. This can be seen in the first triple of the example given in Section 3: (*I; must remind; him the Democrats have controlled the Congress for the last twenty-two years*). One way would be to filter out the triples that do not contain named entities in the subject and object part. This approach would keep only the second triples in the example. One additional step to increase the usefulness of triples would be to apply a coreference analysis, changing pronouns by the objects they refer to. After coreference, (*they; wrote; all the tax bills*) would become (*the Democrats; wrote; all the tax bills*), which is more descriptive.

Another way of improving this approach would be to use post-hoc explanation methods such as integrated gradients [23], to identify which embeddings make the highest contribution to the prediction. This could help identify the triples most relevant to the prediction, giving interpretability to the proposed addition, and further input for a fact-checking system.

5. Future work and conclusion

The HYBRINFOX team is interested in neurosymbolic architectures and our aim is generally to improve performance of Language Models by adding structured information from the texts. This approach has to be adapted to misinformation detection or fact-checking settings. In general, we believe that all tasks that are related to factual claims could benefit from adding structured information into their pipeline in order to increase performance.

The proposed approach uses Language Models such as BERT (for OpenIE) or RoBERTa, but could be upgraded by using most recent advances in Large Language Models such as Mistral or ChatGPT. An LLM prompted with instructions could easily perform a similar pipeline:

1. Extract information triples from the text.
2. Select factual triples.
3. Identify if the factual triples are check-worthy.

This approach could help identify which part of the text contains check-worthy information with better accuracy.

To conclude, the proposed approach, enriching Language Models with a level of structured information, has shown promising results in comparison to the use of Language Models alone on the task of check-worthiness estimation. For check-worthiness, the extraction of factual triples from the text helps classification. However, performance was mixed on non-English texts.

Further analyses need to be conducted with other expert systems to further improve performance. As mentioned in the introduction, the definition of what counts as check-worthy is complex. One approach, which we have not tried, might be to consider as checkworthy first and foremost sentences making objective claims. For that purpose, we may piggyback on the methods used in Task 2 of the CheckThat! Lab [24] dealing with the classification of subjective vs objective sentences. We leave that exploration for further work.

Acknowledgements

We thank two anonymous referees for helpful comments and suggestions on the first version of this report. This work was supported by the programs HYBRINFOX (ANR-21-ASIA-0003), FRONTCOG (ANR-17-EURE-0017), and THEMIS (n°DOS0222794/00 and n° DOS0222795/00). PE thanks Monash University for hosting him during the writing of this paper, in the context of the program PLEXUS (Marie Skłodowska-Curie Action, Horizon Europe Research and Innovation Programme, grant n°101086295).

References

- [1] A. Barrón-Cedeño, F. Alam, J. M. Struß, P. Nakov, T. Chakraborty, T. Elsayed, P. Przybyła, T. Caselli, G. Da San Martino, F. Haouari, C. Li, J. Piskorski, F. Ruggeri, X. Song, R. Suwaileh, Overview of the CLEF-2024 CheckThat! Lab: Check-worthiness, subjectivity, persuasion, roles, authorities and adversarial robustness, in: L. Goeriot, P. Mulhem, G. Quénot, D. Schwab, L. Soulier, G. M. Di Nunzio, P. Galuščáková, A. García Seco de Herrera, G. Faggioli, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Fifteenth International Conference of the CLEF Association (CLEF 2024)*, 2024.
- [2] A. Tchechmedjiev, P. Fafalios, K. Boland, M. Gasquet, M. Zloch, B. Zapilko, S. Dietze, K. Todorov, Claimskg: A knowledge graph of fact-checked claims, in: C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, F. Gandon (Eds.), *The Semantic Web – ISWC 2019*, Springer International Publishing, Cham, 2019, pp. 309–324.
- [3] J. Kim, K.-s. Choi, Unsupervised fact checking by counter-weighted positive and negative evidential paths in a knowledge graph, in: D. Scott, N. Bel, C. Zong (Eds.), *Proceedings of the 28th International Conference on Computational Linguistics*, International Committee on Computational Linguistics, Barcelona, Spain (Online), 2020, pp. 1677–1686. URL: <https://aclanthology.org/2020.coling-main.147>. doi:10.18653/v1/2020.coling-main.147.
- [4] J. Thorne, A. Vlachos, Automated fact checking: Task formulations, methods and future directions, in: E. M. Bender, L. Derczynski, P. Isabelle (Eds.), *Proceedings of the 27th International Conference on Computational Linguistics*, Association for Computational Linguistics, Santa Fe, New Mexico, USA, 2018, pp. 3346–3359. URL: <https://aclanthology.org/C18-1283>.
- [5] N. Hassan, C. Li, M. Tremayne, Detecting check-worthy factual claims in presidential debates, in: *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 1835–1838. URL: <https://doi.org/10.1145/2806416.2806652>. doi:10.1145/2806416.2806652.

- [6] F. Alam, S. Shaar, F. Dalvi, H. Sajjad, A. Nikolov, H. Mubarak, G. Da San Martino, A. Abdelali, N. Durrani, K. Darwish, A. Al-Homaid, W. Zaghouani, T. Caselli, G. Danoe, F. Stolk, B. Bruntink, P. Nakov, Fighting the COVID-19 infodemic: Modeling the perspective of journalists, fact-checkers, social media platforms, policy makers, and the society, in: M.-F. Moens, X. Huang, L. Specia, S. W.-t. Yih (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2021*, Association for Computational Linguistics, Punta Cana, Dominican Republic, 2021, pp. 611–649. URL: <https://aclanthology.org/2021.findings-emnlp.56>. doi:10.18653/v1/2021.findings-emnlp.56.
- [7] F. Arslan, N. Hassan, C. Li, M. Tremayne, A benchmark dataset of check-worthy factual claims, in: *International Conference on Web and Social Media*, 2020. URL: <https://api.semanticscholar.org/CorpusID:216870066>.
- [8] P. Nakov, A. Barrón-Cedeño, T. Elsayed, R. Suwaileh, L. Màrquez, W. Zaghouani, P. Atanasova, S. Kyuchukov, G. Da San Martino, Overview of the CLEF-2018 CheckThat! lab on automatic identification and verification of political claims, in: P. Bellot, C. Trabelsi, J. Mothe, F. Murtagh, J. Y. Nie, L. Soulier, E. SanJuan, L. Cappellato, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Springer International Publishing, Cham, 2018, pp. 372–387.
- [9] F. Rayar, M. Delalandre, V.-H. Le, A large-scale TV video and metadata database for french political content analysis and fact-checking, in: *Proceedings of the 19th International Conference on Content-Based Multimedia Indexing, CBMI '22*, Association for Computing Machinery, New York, NY, USA, 2022, p. 181–185. URL: <https://doi.org/10.1145/3549555.3549557>. doi:10.1145/3549555.3549557.
- [10] B. D. Horne, M. G. Gruppi, NELA-PS: A dataset of pink slime news articles for the study of local news ecosystems, *ArXiv abs/2403.13657* (2024). URL: <https://api.semanticscholar.org/CorpusID:268537274>.
- [11] I. Jaradat, P. Gencheva, A. Barrón-Cedeño, L. Màrquez, P. Nakov, ClaimRank: Detecting check-worthy claims in Arabic and English, in: Y. Liu, T. Paek, M. Patwardhan (Eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, Association for Computational Linguistics, New Orleans, Louisiana, 2018, pp. 26–30. URL: <https://aclanthology.org/N18-5006>. doi:10.18653/v1/N18-5006.
- [12] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, *arXiv:1810.04805 [cs]* (2018).
- [13] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized BERT pretraining approach, *CoRR abs/1907.11692* (2019). URL: <http://arxiv.org/abs/1907.11692>. arXiv:1907.11692.
- [14] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, Q. V. Le, XLNet: generalized autoregressive pretraining for language understanding, Curran Associates Inc., Red Hook, NY, USA, 2019.
- [15] F. Alam, A. Barrón-Cedeño, G. S. Cheema, G. K. Shahi, S. Hakimov, M. Hasanain, C. Li, R. Míguez, H. Mubarak, W. Zaghouani, P. Nakov, Overview of the CLEF-2023 CheckThat! lab Task 1 on check-worthiness in multimodal and multigenre content, in: *Conference and Labs of the Evaluation Forum*, 2023. URL: <https://api.semanticscholar.org/CorpusID:264441760>.
- [16] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, G. Neubig, Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing, *ACM Comput. Surv.* 55 (2023). URL: <https://doi.org/10.1145/3560815>. doi:10.1145/3560815.
- [17] Q. Dong, L. Li, D. Dai, C. Zheng, Z. Wu, B. Chang, X. Sun, J. Xu, L. Li, Z. Sui, A survey for in-context learning, *ArXiv abs/2301.00234* (2023). URL: <https://api.semanticscholar.org/CorpusID:263886074>.
- [18] L. Yufeng, P. Rrubaa, Z. Arkaitz, FactFinders at CheckThat! 2024: Refining check-worthy statement detection with LLMs through data pruning, in: *Working Notes of CLEF 2024 - Conference and Labs of the Evaluation Forum, CLEF '2024*, Grenoble, France, 2024.
- [19] B. Mehmet Eren, K. Kaan Efe, K. Mucahid, TurQUaz at CheckThat! 2024: A hybrid approach of fine-tuning and in-context learning for check-worthiness estimation, in: *Working Notes of CLEF 2024 - Conference and Labs of the Evaluation Forum, CLEF '2024*, Grenoble, France, 2024.
- [20] K. Kolluru, V. Adlakha, S. Aggarwal, p. Mausam, S. Chakrabarti, OpenIE6: Iterative Grid Labeling and Coordination Analysis for Open Information Extraction, in: *The 58th Annual Meeting of the*

Association for Computational Linguistics (ACL), Seattle, U.S.A, 2020.

- [21] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, arXiv preprint arXiv:1607.04606 (2016).
- [22] Y. Ro, Y. Lee, P. Kang, Multi²OIE: Multilingual open information extraction based on multi-head attention with BERT, in: T. Cohn, Y. He, Y. Liu (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2020, Association for Computational Linguistics, Online, 2020, pp. 1107–1117. URL: <https://aclanthology.org/2020.findings-emnlp.99>. doi:10.18653/v1/2020.findings-emnlp.99.
- [23] I. Čík, A. D. Rasamoelina, M. Mach, P. Sinčák, Explaining deep neural network using layer-wise relevance propagation and integrated gradients, in: 2021 IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI), 2021, pp. 000381–000386. doi:10.1109/SAMI50585.2021.9378686.
- [24] J. M. Struß, F. Ruggeri, A. Barrón-Cedeño, F. Alam, D. Dimitrov, A. Galassi, G. Pachov, I. Koychev, P. Nakov, M. Siegel, M. Wiegand, M. Hasanain, R. Suwaileh, W. Zaghouni, Overview of the CLEF-2024 CheckThat! lab task 2 on subjectivity in news articles, 2024.