A Twin XCBR System Using Supportive and Contrastive Explanations

Betül Bayrak[1,*], Kerstin Bach[1]

[1]Norwegian University of Science and Technology (NTNU), Høgskoleringen 1, 7034 Trondheim, Norway

Abstract

Machine learning models are increasingly being applied in safety-critical domains. Therefore, ensuring their trustworthiness and reliability has become a priority. Uncertainty measures the lack of trust in these models, and explanation systems designed as twin systems can provide insights into model decisions to users. Case-based reasoning (CBR) is an experience-based problem-solving methodology with applications across various domains. In this work, we propose a novel approach to generate a twin system, specifically a multi-agent CBR system (MA-CBR system), which utilizes feature attribution-based Explainable Artificial Intelligence (XAI) techniques to explain black-box models in multi-class classification tasks. The proposed approach provides contrastive or supportive instance-based explanations, enabling users to interpret model outputs. Furthermore, we introduce an evaluation metric to assess the system's quality based on its supportiveness for the performance of the underlying black-box model, which we measure through a confidence score. To evaluate the performance of our approach, we apply it to three distinct datasets with differing characteristics. Our results demonstrate the effectiveness of the proposed approach in generating explanations for black-box models in multi-class classification tasks.

Keywords

Explainable Artificial Intelligent (XAI), Explanation Case-Based Reasoning (XCBR), Model-Agnostic Explanation Generation, Twin XAI Systems

1. Introduction

The relationship between the performance and interpretability of machine learning models, which are often considered to have a trade-off between complexity and interpretability, has been a topic of discussion. A machine-learning model with low interpretability and high opacity can be called a black-box model [1]. Due to the low level of understandability by individuals in the applications of black-box models, especially as their real-world use increases, interest has started to grow in XAI systems that help understand the reasons behind the decisions made by these models. As black-box models are increasingly applied in safety-critical domains, ensuring their trustworthiness and reliability has become a priority, and uncertainty is defined as a measure of the lack of trust in these models [2]. Considering the demand for enhancing

ICCBR XCBR'23: Workshop on Case-Based Reasoning for the Explanation of Intelligent Systems at ICCBR2023, July 17 – 20, 2023, Aberdeen, Scotland

*Corresponding author.
betul.bayrak@ntnu.no (B. Bayrak); kerstin.bach@ntnu.no (K. Bach)
https://www.ntnu.edu/employees/betul.bayrak/ (B. Bayrak); https://www.ntnu.edu/employees/kerstin.bach/ (K. Bach)
0000-0002-0554-9823 (B. Bayrak); 0000-0002-4256-7676 (K. Bach)

understandability and ensuring reliability, XAI systems designed as twin systems are suitable to meet the requirements for providing insights into model decisions to users. Further information about twin-XAI systems can be found in Section 2.2.

In XAI systems, the dimensions of explanations can be simplified as global and local explanations. Global explanations offer insights into understanding the overall logic of a model and encompass the entire reasoning process that leads to all the different possible outcomes. On the other hand, local explanations focus on understanding the specific reasons behind a particular decision, such as a single prediction or decision. [1]

Case-based reasoning (CBR) is a problem-solving approach that utilizes past experiences and has broad applicability in various domains. By leveraging the flexibility and interpretability of the CBR methodology, eXplanation CBR (XCBR) systems, CBR systems designed to explain a model, can provide global and/or local explanations. Additionally, these systems are adaptable to changes in data distribution and can generate trustworthy explanations using a small amount of data [3].

This work proposes a novel approach to generate a twin XAI system. The twin system is developed as a multi-agent CBR (MA-CBR) system that utilizes feature attributions to explain multi-class classification black-box models. In the XCBR system, an agent is developed for each class, and every agent is modeled separately. In the modeling phase, feature attributions and data distribution are used. Thereby, this approach projects the different characteristics of classes through feature attribution. The proposed approach provides contrastive or supportive instance-based explanations that enable users to interpret model outputs. Furthermore, an evaluation metric is introduced to assess the adaptability of the system based on supportiveness for the performance of the underlying black-box model, which is $rigidity$.

To evaluate the performance of our approach, we apply it to three distinct datasets with differing characteristics. Our results demonstrate the effectiveness of the proposed approach in generating interpretable explanations for black-box models in multi-class classification tasks.

The main contributions of this work are as follows:

- Facilitate the incorporation of expert knowledge into the XCBR system, thereby improving the reliability and trustworthiness of the explanations provided.
- Multi-agent structure enables the generation of instance-based explanations that incorporate locality and globality in the explanations.
- Proposes an evaluation metric, $rigidity$, to measure the adaptability of the black-box model's performance through the proposed explanation system.
- Provides reproducible benchmarking experiments and open-source implementation of the proposed approach and evaluation metric (https://github.com/b-bayrak/Twin_XAI).

The rest of the paper is structured as follows. Section 2 provides an overview of background information and related work. In Section 3, details of the proposed approach and evaluation metric are described. Conducted experiments, use cases, results, discussions, and future work directions are given in Section 4. Finally, Section 5 presents a brief conclusion of the paper.

2. Background and Related Work

This paper combines the following lines of related research: Uncertainty focused XAI, Twin XAI systems, and CBR for XAI.

2.1. Uncertainty

Uncertainty in black-box models has been an active discussion topic of research in machine learning and artificial intelligence. Several techniques have been proposed to measure and represent uncertainty in black-box models, such as Monte Carlo dropout [4] and Bayesian network's weight patterns [5]. However, these techniques often suffer from high computational costs and limited applicability to different data types. To address this issue, decision-support XAI systems have been proposed to provide users with a better understanding of the uncertainty inherent in black-box models. These systems typically generate explanations that highlight the key features that contribute to the model's output and provide a measure of the confidence or uncertainty associated with each prediction [6, 5, 7]. Overall, using XAI systems to represent uncertainty in black-box models has shown promise in improving user trust and understanding of the model's output.

2.2. Twin Explanation Systems

Considering the demand for enhancing understandability and ensuring reliability, twin XAI systems are suitable to meet the requirements for providing insights into model decisions to users.

Twin XAI systems consist of two separate models trained on the same dataset. One of the models acts as the primary model, the black-box model, and the explainer model provides explanations for black-box model decisions. The explainer model provides insights into how the black-box model makes its predictions, thereby providing greater transparency and accountability, and it can be built as different types like a machine learning algorithm[8], rule-based model [9], or CBR system [10].

2.3. CBR Methodology for XAI Systems

CBR has been applied to various domains, including medical diagnosis [11], and financial fraud detection [12]. Recently, there has been an emerging interest in developing twin XCBR systems, which generate instance-based explanations that allow incorporating the effects of both local and global features. For instance, Bayrak et al. [13] proposed a twin XCBR system that utilizes feature attributions to explain multi-class classification black-box models, and Ahmed et al. [14] showed that a CBR model as an explainer can perform good enough with the help of the additive models.

3. Twin XCBR System

In this section, we comprehensively describe the methodology utilized to construct the proposed twin XCBR system.

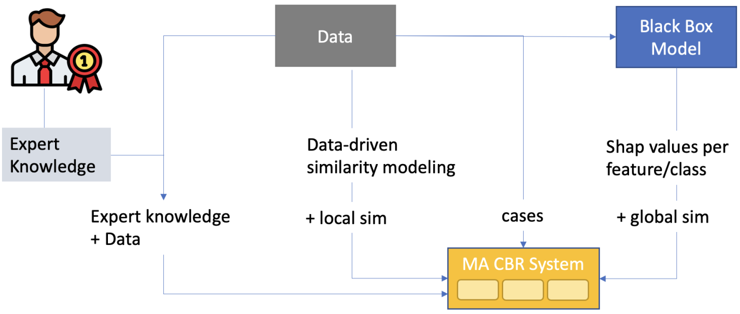

Figure 1: Components and their intended use to develop proposed MA-CBR system.

3.1. Multi-agent Structure

To build the MA-CBR system, which functions as an explainer system as depicted in Figure 1, both a dataset and a black-box model are required. The input data set may be the same as the one utilized for training the black-box model, a different data set with a similar concept and distribution, or a combination of both. In all circumstances, the input samples must be labeled since all of the samples need to be grouped by their labels (i.e., classes), and each distinct class necessitates the creation of an agent for the MA-CBR system design. Each agent has a case base for which the global and local similarity measures are developed independently. Consequently, this approach allows for the projection of diverse characteristics of distinct classes through feature attributions and expert knowledge.

SHAP values calculated separately for each class over the black-box model are employed as weights for the CBR agent's global similarity measure. To develop the local similarity measures, a data-driven similarity measure development method is employed [15]. In this approach, Verma et al. proposed an Inter Quartile Range-based polynomial modeling. For both global and local similarity measure developments, expert knowledge (if available) may be incorporated (refer to Section 3.2).

Following the development of similarity measures for each case base, data samples (cases) are classified according to their respective labels and subsequently incorporated into the appropriate case base. After completing these procedures, the MA-CBR system is deemed ready for querying.

3.2. Injecting Expert Knowledge

Under circumstances where expert knowledge is available, different types of domain knowledge provided by experts can be incorporated into the system. The incorporation of domain knowledge can be done on two different levels:

- At the *data level*, it can be integrated into the MA-CBR system in various ways. This includes the exclusion of existing features, the inclusion of new features, or the combination and compression of existing features. Additionally, it can be utilized to eliminate invalid cases using rule-based domain-specific conditions.

- At the *similarity measurement level*, it can be incorporated into both global and local similarity measurements. In the global similarity, domain knowledge can be used to improve feature importance weights that are calculated by SHAP values and their effects on different categories. In the local similarity, domain knowledge can help to define relationships between different categories and the ranges of features.

Domain knowledge can be used during the system design process or can be added retroactively, even after the system has already been put into use. The system's flexible structure allows for such additions or modifications.

Also, expert knowledge can provide additional insights and enable more transparent and interpretable decision-making processes, thereby incorporating expert knowledge into the explanation system plays a crucial role in ensuring compliance with the General Data Protection Regulation (GDPR) in terms of addressing concerns around bias and meeting regulatory requirements related to data protection and privacy.

3.3. Explanation Generation and Representation

After building up the MA-CBR system, this section represents how the explanations are generated and represented. As shown in Figure 2, the system's input is an explanation case, e_i. An explanation case consists of a data sample x_i and its prediction result y_i generated by the black-box model. Where $x_i = a_0, a_1, ..., a_m$, $y_i = BBM(x_i)$ and $e_i = x_i, y_i$.

When a new explanation case, e_i, arrives at the explainer, a new query is conducted by all agents. Each agent returns the most similar cases with their corresponding similarity scores. The query result with the highest similarity score, e_w, is selected as the explanation source and used for comparison. The similarity score of the selected query result is used as the confidence score for the explanation. A class comparison is made between the selected explanation source and e_i to see if they are identical. If they are identical, this indicates agreement between the CBR system's result and the black-box model. The winning query result, e_w, will be used as a supportive explanation with a confidence score will be provided. If they are not the same, e_w as a contrastive explanation with a confidence score will be provided.

In both circumstances, the proposed system provides a clear and informative output, consisting of the explanation case, e_i, the winning query result, e_w, and a confidence score that reflects the level of agreement or disagreement with the black-box model. Moreover, the system presents a sorted version of attributes based on feature importance and difference. The inclusion of both e_i and e_w, with the attribute order, allows users to reason semantically between them and gain a comprehensive understanding of the prediction process.

Considering the demand for enhancing understandability and ensuring reliability, the proposed supportive/contrastive explanations are suitable to meet the requirements for providing insights into model decisions to users.

3.4. Evaluation Technique

Assuming that X_{test} contains n data samples, $X_{test} = x_1, ..., x_n$, and y represents the set of prediction results made by black-box model, where $y = y_1, ..., y_n$ and $y_i = BBM(x_i)$. And, S is a subset of the X_{test} consisting of supported predictions by the explainer. The support score,

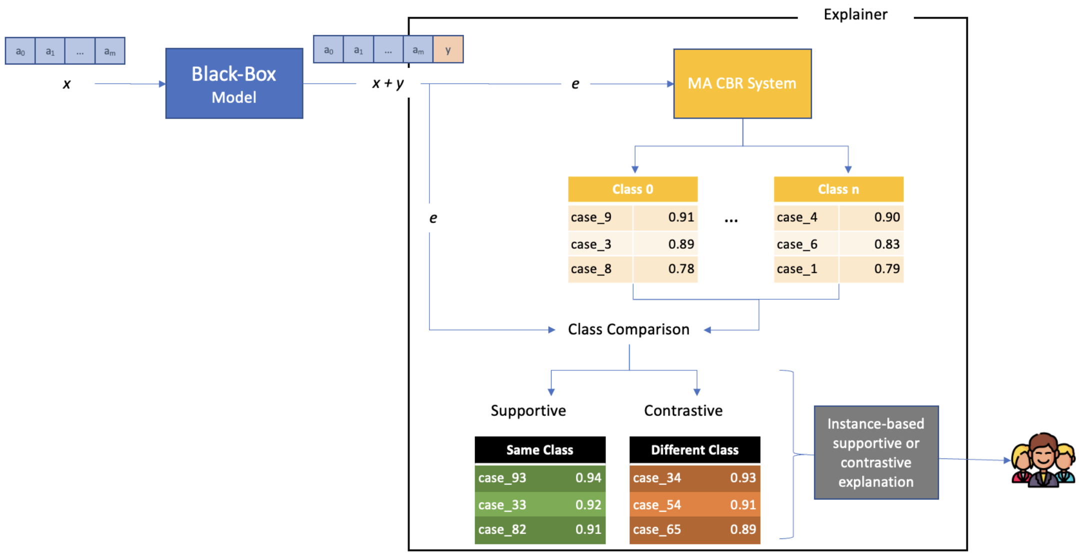

Figure 2: Explanation generation process using the proposed MA-CBR system.

$support$, measures the ratio of the supported black-box model decisions by the explanation system.

$$support = \frac{|S|}{n} \tag{1}$$

The expected behavior of the explainer for the $support$ score is to be approximately proportional to black-box model's accuracy. Here, acc denotes the accuracy score of black-box modelcalculated as the comparison between y and ground truth. For example, if the black-box modelis a high performance model, a high $support$ score is expected because the idea of the explanation is to support the black-box modelprediction with a supportive instance when it is correct.

The proposed evaluation technique, $rigidity$, measures the adaptability of the explainer to the black-box model's performance.

$$rigidity = |1 - \frac{support}{acc}| \tag{2}$$

A low $rigidity$ indicates better performance by the explainer with the black-box model. The lower the value, the better.

4. Experiments

In order to ensure that our approach is applicable to different domains and evaluate the effectiveness of our proposed approach, we conducted experiments on three datasets which were carefully selected to represent a variety of characteristics.

4.1. Use Cases

In the use cases, a standardized approach was employed. The datasets were partitioned into training and testing sets to ensure a balanced representation across classes. The cases were then constructed using the training data and imported into the relevant case bases. Local similarity measures were set using the calculated IQRs of attributes as explained in Section 3.1 and the global similarity measures of the case bases were set using calculated SHAP values for each class separately. With the explainer system in place, explanation cases were provided as inputs, and supportive or contrastive explanations were generated as shown in Fig 1. This systematic process ensured consistency and facilitated the evaluation of the proposed approach.

4.1.1. *Use-case 1*: Depression Screening Dataset

The dataset[16] used in this use-case was collected to measure the level of depression among undergraduate students. The participants consisted of undergraduate students from Tecnológico Nacional de México (TecNM)/Instituto Tecnológico de Mérida (ITM) between May 2020 and December 2020, ranging in age from 17 to 23 years. All relevant guidelines and regulations were adhered to, and the students provided their consent to participate in the study. As part of the study, the students were required to complete a 102-item questionnaire, where a response of "true" was assigned a value of 1 and "false" a value of 0.

This use case has been presented in a preliminary study of this paper. The dataset and domain knowledge was provided as part of the XCBR challenge track at the 2022 International Conference on Case-Based Reasoning (ICCBR-2022), and details of the preliminary experiments can be found in [13].

The dataset consists of 105 samples and is designed for a 3-class classification model of depression levels, with the three possible classes indicating severity levels with varying numbers of samples per class. Domain knowledge is provided as "expected answers" that represent the expected responses from an individual with depression. Domain knowledge is incorporated into the explanation system as a new attribute called "Matched", which represents the number of overlapping items between actual and expected answers. For the experiments, 20% of the dataset is used as test data, while the remaining 80% (train data) and its oversampled version are used to train the black-box models and build the explanation systems.

We conducted three different experiments for this use case. The first experiment compared the performance of CBR systems with that of black-box models using raw and oversampled data and the second experiment compared the performance of CBR systems with and without domain knowledge. In the third experiment, cases were constructed using the train data and the calculated "Matched" attribute (domain knowledge). With the first and second experiments, the performance of the CBR system improved from 0.29 to 0.52 accuracy score using raw data with domain knowledge. In the third experiment, the global similarity measures of the case bases were set using calculated SHAP values for each class separately for the MA-CBR system. The cases were constructed using the train data and the calculated "Matched" attribute (domain knowledge). In this experiment, the test set comprised 21 instances, and the built black-box model, an MLP, achieved an accuracy score of 0.24. Explanations were generated for each instance in the test set, and the system supported 10 decisions out of the 21 decisions made by the black-box model. The *rigidity* of the explanation system was calculated to be 0.984.

4.1.2. *Use-case 2*: SelfBack App Usage Prediction Dataset

SelfBack project[1] developed a decision support system to improve self-management of non-specific low back pain [17]. As part of the project, a mobile application was developed to convey personalized self-management content through exercises and educational content. The App Usage dataset is derived from the collected data, and in this use case, we predict the usage of the mobile application given the answers to a baseline questionnaire that characterizes a user's situation with regard to their current episode of back pain. The underlying question in this use case is whether the suggested intervention would suit a user. Identifying patients with lower or moderate app engagement enables personalized reminders, interactive exercises, and educational content to engage them actively. These interventions aim to enhance patient involvement, treatment adherence, and overall outcomes.

The App Usage dataset contains data from 230 users and 26 continuous and nominal features derived from users' input (baseline questionnaire). Data instances were labeled using 3 classes indicating levels of app usage with varying numbers of samples per class. This use-case applies real-world data, and domain knowledge was incorporated to model the local similarity measures. For the experiments, 30% of the dataset was used as test data, while the remaining 70% (train data) was used to train the black-box models and build the explanation systems.

The test set comprised 69 instances, and the built black-box model(an MLP) achieved an accuracy score of 0.58. Explanations were generated for each instance in the test set, and the system supported 48 decisions out of the 69 decisions made by the black-box model. The *rigidity* of the explanation system was calculated to be 0.199.

4.1.3. *Use-case 3*: Wine Quality Dataset

The dataset used in this use-case is related to red variants of the Portuguese "Vinho Verde" wine [18]. It contains 4898 rows and 12 features, with ordered and unbalanced classes. There is no domain knowledge incorporated in this use-case. For the experiments, 30% of the dataset is used as test data, while the remaining 70% (train data) is used to train the black-box models and build the explanation systems.

The test set consisted of 1950 instances, and for this use-case, we trained five models: *K Nearest Neighbors Classifier*, *MLP*, *Decision Tree Classifier*, *Gradient Boosting Classifier*, and *Random Forest Classifier*. They achieved accuracy scores of 0.749, 0.767, 0.776, 0.839, and 0.848 respectively. The same procedure was applied to all models using the same train-test set split (refer to Figure 3).

To compare with other use-cases, we selected the MLP model, which achieved an accuracy score of approximately 0.77. The system supported 1930 decisions out of the 1950 decisions made by the black-box model. The *rigidity* of the explanation system was calculated to be 0.2197.

[1] https://www.selfback.eu/

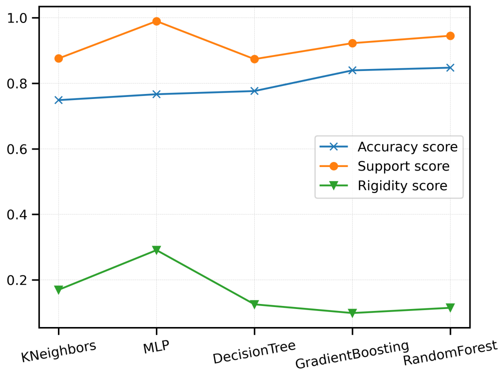

Figure 3: Accuracy, Support, and Rigidity scores of different models over Wine Quality Dataset.

Table 1
Performance evaluation of the proposed approach on three use cases with accuracy, support, and rigidity scores.

	accuracy	support	rigidity
Use-case 1	0.24	0.476	0.984
Use-case 2	0.58	0.696	0.199
Use-case 3	0.77	0.9897	0.2197

4.2. Discussion

This paper aims to propose a novel approach to generate a twin XAI system and apply the approach to diverse datasets to assess the robustness and generalizability across different domains. In the presented use cases, MLP models were trained as black-box models and used to build the twin systems without refining them. Some of the models performed with very low accuracy, while others had comparably better accuracy scores (see Table 1). Also in Figure 3, *use-case 3*, the same setup is used with 5 models with different performances. These experiments allowed us to demonstrate the applicability of the system. The approach is applied to two data sets with a very limited number of samples (*use-case 1 and 2*) and one dataset with a large number of samples (*use-case 3*), and all data sets are unbalanced. In *use-case 1*, the first experiment showed that the built CBR models perform comparably better than the built black-box models. As a result, we showed that the proposed approach is applicable to unbalanced data sets and has considerable performance with data sets of different sizes. This is an essential feature in real-world explanation systems.

As mentioned in Section 3.2, expert knowledge can be incorporated into the system in different ways. For example, in *use-case 1*, the provided expert knowledge is incorporated as a new attribute. In *use-case 2*, it was used to model similarities, while no expert knowledge was used in *use-case 3*. In either case, the explanation system was built successfully. The contribution of domain knowledge is clearly shown in the second experiment of *use-case 2*, see Section 4.1.2.

As stated above, the wine quality (*use-case 3*), app usage(*use-case 2*), and depression screening (*use-case 1*) applications are built on a black-box modelwith 0.77, 0.58, and 0.24 accuracy scores, respectively. Also, their *support* scores are calculated as 0.9897, 0.696, and 0.476, and *rigidity* scores are 0.2197, 0.199, and 0.984, respectively. As can be understood from the results, when the accuracy score decreases, the *support* score also decreases, as expected in an adaptable explanation system. A low *rigidity* score indicates better performance, so in a perfect explanation system, the *rigidity* score would be 0. In our experiments, similar to the expected behavior of the explainer, the *rigidity* scores are approximately proportional to black-box models' accuracies. Also, in *use-case 3* with 0.77 accuracy score and 0.9897 support score performs worse than the other two in terms of *rigidity* because 1930 of the 1950 decisions made by the black-box modelare supported, and the performance of the black-box modelis not perfect, so a lower number of supported decisions is expected.

As mentioned before, we expect a flexible system's rigidity score to approach 0. In *use-case 3*, we anticipated an approximately linear relationship between accuracy and support scores. Figure 3 demonstrates that, in most cases, as the accuracy score increases, the support score also increases accordingly, with the exception of the MLP model. In the case of the MLP model, the support score approaches 1 while the accuracy score is around 0.77. This observation suggests that the MLP model and the constructed MA-CBR systemexplainer exhibit similar behaviors, indicating a certain level of dependence on the explainer's performance. This dependency presents a challenge that should be addressed in future work.

In the use cases, data sets with different data types and characteristics were used. However, due to the nature of the tools used, only tabular data was used. For future work, this approach can be extended to different kinds of data, such as images, using a similar approach to Barnett et al.'s work [19]. Another area for improvement is the representation of the explanations and the measurement of the effect of the explanations on the users through a user study.

5. Conclusion

In conclusion, we have proposed a novel approach to generate a twin XAI system that utilizes feature attributions to explain multi-class classification black-box models. Our approach employs a MA-CBR system, where an agent is developed for each class, modeling their similarity measures separately. By projecting the different characteristics of classes through feature attribution, our approach provides contrastive or supportive instance-based explanations that enable users to interpret model outputs. Moreover, we introduced an evaluation metric, rigidity, to assess the system's quality based on supportiveness for the performance of the underlying black-box model. Through experiments on three distinct datasets with differing characteristics, we demonstrated the effectiveness and applicability of our approach in generating explanations for black-box models in multi-class classification tasks. Our work also facilitates the incorporation of expert knowledge into the XCBR system, improving the reliability and trustworthiness of the explanations provided, and provides a reproducible benchmarking experiment and open-source implementation of the proposed approach and evaluation metric. While our approach currently only supports tabular data, it can be extended to other data types, such as images, using a similar approach to previous works. Our explanation system can be useful in various domains, including healthcare, finance, and law, where XAI is essential. Our proposed approach contributes to the growing research on explainable AI and can provide valuable insights for stakeholders in various domains.

References

[1] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, ACM computing surveys (CSUR) 51 (2018) 1–42.

[2] K. E. Brown, D. A. Talbert, Using explainable ai to measure feature contribution to uncertainty, in: The International FLAIRS Conference Proceedings, volume 35, 2022.

[3] B. Bayrak, K. Bach, When to Explain? Model Agnostic Explanation Using a Case-based Approach and Counterfactuals, in: A. Rutle (Ed.), Norsk IKT-konferanse for forskning og utdanning, 1, 2022.

[4] G. Lee, Monte-carlo dropout based uncertainty analysis in input attributions of multivariate temporal neural networks (2020).

[5] K. Bykov, M. M.-C. Höhne, K.-R. Müller, S. Nakajima, M. Kloft, How much can i trust you?–quantifying uncertainties in explaining neural networks, arXiv preprint arXiv:2006.09000 (2020).

[6] A. Dubey, N. Lyons, A. Santra, A. Pandey, Xai-bayeshar: A novel framework for human activity recognition with integrated uncertainty and shapely values, arXiv preprint arXiv:2211.03451 (2022).

[7] J. Jiang, S. Kahai, M. Yang, Who needs explanation and when? juggling explainable ai and user epistemic uncertainty, International Journal of Human-Computer Studies 165 (2022) 102839.

[8] F. Silva-Aravena, H. Núñez Delafuente, J. H. Gutiérrez-Bahamondes, J. Morales, A hybrid algorithm of ml and xai to prevent breast cancer: A strategy to support decision making, Cancers 15 (2023) 2443.

[9] D. Macha, M. Kozielski, Ł. Wróbel, M. Sikora, Rulexai—a package for rule-based explanations of machine learning model, SoftwareX 20 (2022) 101209.

[10] E. M. Kenny, M. T. Keane, Explaining deep learning using examples: Optimal feature weighting methods for twin systems using post-hoc, explanation-by-example in xai, Knowledge-Based Systems 233 (2021) 107530.

[11] S. Sharma, D. Mehrotra, Two-stage cbr based healthcare model to diagnose liver disease, International Journal of Computing and Digital Systems 10 (2021) 1–8.

[12] I. Eweoya, A. Adebiyi, A. Azeta, A. E. Azeta, Fraud prediction in bank loan administration using decision tree, in: Journal of Physics: Conference Series, volume 1299, IOP Publishing, 2019, p. 012037.

[13] B. Bayrak, P. Marin Veites, K. Bach, Explaining your neighbourhood: A CBR approach for explaining black-box models (2022) 251–255. URL: http://ceur-ws.org/Vol-994.

[14] M. U. Ahmed, S. Barua, S. Begum, M. R. Islam, R. O. Weber, When a cbr in hand is better than twins in the bush, arXiv preprint arXiv:2305.05111 (2023).

[15] D. Verma, K. Bach, P. J. Mork, Similarity measure development for case-based reasoning–a data-driven approach, in: Symposium of the Norwegian AI Society, Springer, 2019, pp. 143–148.

[16] M. G. Orozco-del Castillo, E. C. Orozco-del Castillo, E. Brito-Borges, C. Bermejo-Sabbagh, N. Cuevas-Cuevas, An artificial neural network for depression screening and questionnaire refinement in undergraduate students, in: M. F. Mata-Rivera, R. Zagal-Flores (Eds.), Telematics and Computing, Springer International Publishing, Cham, 2021, pp. 1–13.
[17] L. F. Sandal, K. Bach, C. K. Øverås, M. J. Svendsen, T. Dalager, J. S. D. Jensen, A. Kongsvold, A. L. Nordstoga, E. M. Bardal, I. Ashikhmin, et al., Effectiveness of app-delivered, tailored self-management support for adults with lower back pain–related disability: a selfback randomized clinical trial, JAMA internal medicine 181 (2021) 1288–1296.
[18] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, J. Reis, Modeling wine preferences by data mining from physicochemical properties, Decision support systems 47 (2009) 547–553.
[19] A. J. Barnett, F. R. Schwartz, C. Tao, C. Chen, Y. Ren, J. Y. Lo, C. Rudin, A case-based interpretable deep learning model for classification of mass lesions in digital mammography, Nature Machine Intelligence 3 (2021) 1061–1070.