

# ModelSpeX: Model Specification Using Explainable Artificial Intelligence Methods

Udo Schlegel<sup>1,2</sup>, Eren Cakmak<sup>1</sup>, and Daniel A. Keim<sup>1</sup>

<sup>1</sup> University of Konstanz, Germany    <sup>2</sup> Untangle AI, Singapore

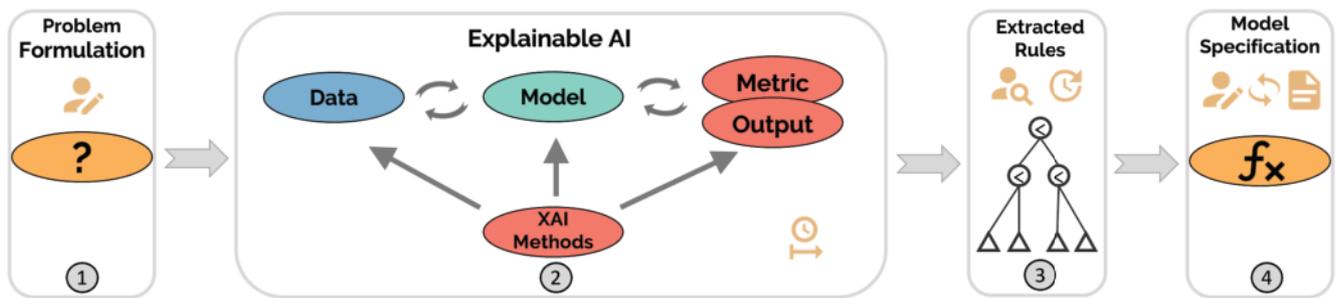


Figure 1: The ModelSpeX workflow: (1) Begins with the user task (problem formulation); (2) The user extracts semi-automatically rules using visual analytics interfaces and XAI methods for rule extraction. (3) The user iteratively refines the rule sets to generate a model specification, which (4) enables to evaluate and verify trained models by comparing them to the formulated problem.

## Abstract

*Explainable artificial intelligence (XAI) methods aim to reveal the non-transparent decision-making mechanisms of black-box models. The evaluation of insight generated by such XAI methods remains challenging as the applied techniques depend on many factors (e.g., parameters and human interpretation). We propose ModelSpeX, a visual analytics workflow to interactively extract human-centered rule-sets to generate model specifications from black-box models (e.g., neural networks). The workflow enables to reason about the underlying problem, to extract decision rule sets, and to evaluate the suitability of the model for a particular task. An exemplary usage scenario walks an analyst through the steps of the workflow to show the applicability.*

## 1. Introduction

Artificial intelligence (AI) efficiently solves previously challenging problems (e.g., machine translation [WSC\*16], autonomous driving [CSKX15]) due to recently developed machine learning (ML) methods. These ML methods are, however, often criticized for being black-box models (e.g., neural networks) with non-transparent decision-making, impractical for critical domains (e.g., medicine) [Rud18] leading to the development of explainable AI (XAI) [GMR\*18]. The primary goals of XAI are to reveal the underlying non-transparent decision-making mechanisms of black-box models [AB18], to understand, to debug, to optimize models [LWLZ17], to build trust [GMR\*18], to increase transparency, fairness, as well as privacy [DK17]. Explanations extracted by XAI methods (e.g., LIME [RSG16]) require human interpretation, which can lead to potentially wrong insights [RTC18]. Such explanations, therefore, require additional context information and

new methods to quantify as well as verify the extracted insight about the underlying decision-making process [HMKL18]. For instance, feature importance or attributions (e.g., LIME [RSG16]) can be misleading (biased) due to unbalanced data (e.g., showing only the describing feature for the majority class) [SHJ\*19]. As a result of such biases, the number of explanations required to unveil the internal decision-making processes of a black-box model, e.g., deep neural networks, remains unanswered [MZR18]. The human inspection of complicated decision rules sets in the context of the data, model, and the underlying problem allows to adjust (e.g., prune) and to evaluate the extracted explanations and knowledge. For example, decision rule extraction methods produce comprehensible explanations (rules) for humans, which help to reveal the decision-making process and potential biases [Hai16]. To evaluate and verify the quality of generated explanations, we propose a workflow to interactively create machine-readable descriptions that aim to summarize the set of tasks trained models to solve.

We propose ModelSpeX, a workflow that uses XAI methods (e.g., decision rule extraction techniques) to generate an understandable model specification (a machine-readable description) for comparison, evaluation, and verification of complex models such as deep neural networks. The main goal of the proposed workflow is to combine rule extraction XAI methods with interactive visualizations to highlight the decision-making process of black-box models and to generate descriptive model specifications. Such model specifications allow us to evaluate models and verify the suitability to the underlying problem enabling humans to compare the trained model with the underlying task (problem). Overall, the model specification aims to divide complex formulated problems into sub-problems to enable domain experts to find transparent and straightforward solutions for the original problem. In summary, the main contributions of this paper are: (1) A conceptual workflow to extract model specifications from XAI methods to facilitate human users to understand and assess whether a model is suitable for particular tasks. (2) An exemplary real-world use case to highlight suitable XAI techniques for our workflow.

## 2. Related Approaches

Depending on the research field, the definitions for the term explainability differ, for example, in machine learning [Lip18] and visual analytics [SSSE19]. However, most definitions have similar motivations, such as fairness, privacy, reliability, and trust building [DK17, GMR\*18]. The post-hoc application of XAI methods onto black-box models (after training steps) produces explanations for predictions [AB18]. In recent years, the number of XAI methods increased heavily intending to achieve explainability, especially in computer vision, helping XAI to advance [SAE\*19].

Workflows conceptualize the general task of discovering and extracting knowledge from data. Prominent workflows are, for instance, the knowledge discovery in databases (KDD) process [PF91] and the visual analytics model [KAF\*08]. Further, Sacha et al. [SSZ\*17] introduced interactions for the various steps of a machine learning pipeline to supplement an analyst to, e.g. support direct model manipulation. However, many of the proposed interactions cannot be applied as the training of deep neural networks are often time-consuming (days to weeks). Focusing on exploratory model analysis, Cashman et al. [CHH\*19] present a workflow to incorporate visual analytics principles in the model selection process. The workflow, however, lacks methods to investigate the behavior of black-box models. Spinner et al. [SSSE19] proposed a framework and pipeline for XAI to incorporate various concepts, such as understanding, diagnosis, and refinement [LWLZ17]. However, the proposed framework does not describe the evaluation and verification of explanations in the context of user tasks at hand (problem).

In summary, many visual analytics workflows lack the extension to explain decisions of black-box models in a comprehensive way. Most workflows miss a user-centered evaluation of XAI explanations in the context of the problem formulation. We, therefore, propose a workflow that incorporates decision rule explanations for specifying as well as evaluating models and their underlying formulated problems (user tasks).

## 3. ModelSpeX Workflow for Model Specifications

We propose a ModelSpeX, a model specification workflow that describes the required steps to compare and assess the relationships between the problem (task) and the learned model approximations. The workflow (see Fig. 1) consists of four steps and ties together visual analytics with XAI methods to extract human-centered explanations and evaluate the black-box solution. The overall goal of the workflow is to interactively generate a model specification to reason about the underlying user-defined task (problem). We define a model specification as a machine-readable description, for instance, a decision rule set that describes a solution for a problem. A model specification can be used to

- evaluate models (e.g., comparison to formulated problems);
- refine problem formulations (e.g., add new classes);
- generate new knowledge about a specific application domain (e.g., identification of decision boundaries around certain data samples);
- and to interactively align the user task with the model (e.g., identify the model that meets the expectations of a domain expert).

In the following section, we will further detail the steps and interactions involved to generate a model specification on a use case and discuss potential methods for each step using the real-world example of a recurrent neural network (RNN) for anomaly detection in time series sensor data for predictive maintenance [Mob02].

### 3.1. Generate Problem Formulation

Machine learning starts with an analytical task (formulated problem) about a dataset in a particular application domain (problem domain) [PF91]. The analysis of data involves multiple steps and can be among others described by Norman's Stages of Action cycle [Nor88]. The first step in the execution phase of the model is a problem formulation, that aims to form an intention to solve tasks. A correct problem formulation is a critical part of every ML approach as a faulty formulation results in flawed models [Lip18]. Therefore, Lipton [Lip18] pleads for new research aiming to solve the issue of general problem formulation for machine learning models. In general, such problem formulations can be generated by a domain expert through a visual exploration of the underlying data [CHH\*19]. In contrast to earlier work (e.g., Cashman et al. [CHH\*19]), we propose to interactively generate a model specification for selecting and evaluating relevant black-box models in the visual analytics process.

In our *anomaly detection use case*, the problem is formulated (Fig. 2 (1)) twofold regarding the questions of the analyst:

- Is the anomaly detection (RNN) robust enough for deployment?
- Does the RNN model match domain knowledge?

The first task introduces our problem of anomaly detection with a complex model (RNN) and is answerable with a metric score comparison to a baseline and a first test run on real-world data. However, the complex model could learn some proxy to improve performance. Thus, the second task evaluates model predictions based on the inner workings to foster understanding and trust-building for maintenance engineers.

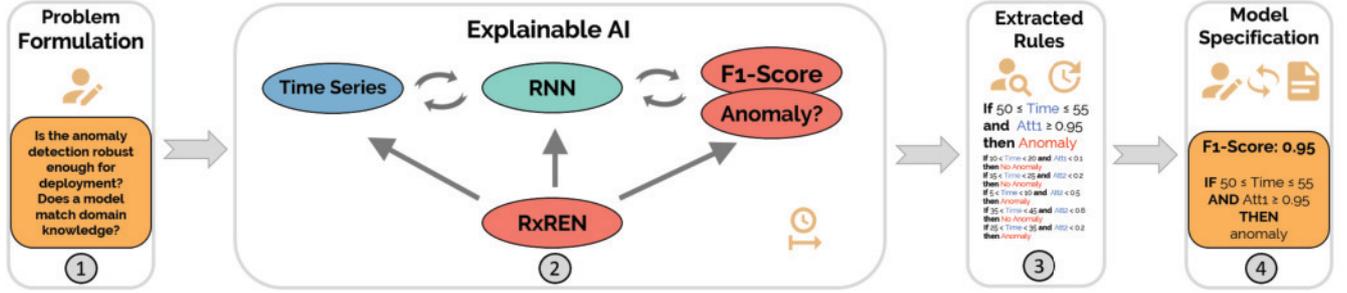


Figure 2: The ModelSpeX workflow applied on an example anomaly detection use case: The formulated problem (1) consists of building trust for domain experts. The model (2) consists of an RNN with accuracy as metric and RxREN as XAI method. The extracted rules (3) are basic if-then clauses to prune. The model specification (4) consists of the model performance and a rule that represents domain knowledge.

### 3.2. Applying Explainable AI

A visual analytics system ideally provides feedback for users and incorporates the analyst in the generation, debugging, and refinement of models [LWLZ17]. The execution phase of Norman’s stages of interaction [Nor88] was used by Sacha et al. [SSZ\*17] to describe a human-centered model generation using interactive visualization. In contrast, our workflow describes the evaluation phase of Norman’s stages of interaction for black-box models using XAI methods and helps to compare and evaluate the specification of the model with the initial problem. As a foundation of our workflow, we extend visual analytics model by Keim et al. [KAF\*08]. We enhance the model by applying XAI methods to the data, model, and output of the model to reveal learned relationships and structures in the underlying dataset to generate rule sets. The extraction of many rules as possible requires to iteratively apply the XAI methods on multiple non-deterministic trained black-box models, which allows understanding the individual local optimum of learned models. The successive rule set extraction also enables to extract temporary learned rules from the training process to reveal critical memorizing moments for the model (e.g., overfitting, learning biases of the data). In the following, we describe explainable AI methods to extract rules sets from the data, model, and output of the model.

The goal of information visualization is to present data to gain knowledge about the real world [BBSS03]. For this purpose, often, models are trained to approximate the underlying problem (e.g., classification) and to identify patterns in the data (e.g., decision boundaries) [KKEM10]. The incremental training process (supervised, semi-supervised, or unsupervised) of such models aims to minimize a metric and capture the complexity of the formulated real-world problem [PF91]. The training of such models poses many challenges due to data bias, unbalanced datasets, overfitting, underfitting, etc. The output of a model (predictions) allows an analyst to assess the performance of a trained model (e.g., accuracy, recall). For instance, examining the output and the loss over time may enable the identification of overfitting as well as underfitting. Decision rules can be extracted from the data, the model, and the output of the model. However, the consideration of only one evaluation metric is not sufficient as there are specific trade-offs between them (e.g., accuracy is sensitive to imbalanced datasets).

**XAI Methods** – In general, extracting decision rules from data can be achieved by decision rule list algorithms, for instance, interpretable decision rule sets [LBL16]. However, such algorithms

have NP-hard complexities [Gol96] and often overfit fast without being able to generalize on large datasets [Hai16]. Post-hoc XAI methods can use a trained model to extract rules based on the inner workings. For instance, DeepRED [ZMJ16] creates decision rules based on the weights and hidden states of neurons in complex neural networks. Such rule extraction algorithms for models often do not incorporate data, and thus are often hard to interpret [Hai16]. Furthermore, the analysis of the output can allow the model to be discarded at an early stage or to be optimized by steering the model into another direction. TREPAN [CS96], for instance, enables to extract decision trees on the basis of sampling the output of a multi-layer perceptron model. As such extraction algorithms work on the data, they are easier to interpret. RxREN [GAK12] improves a sampling approach like TREPAN [CS96] with a decomposition of neuron weights to extract better decision rules on the cost of complexity. However, such algorithms have a high complexity [Gol96] due to sampling input data and use the underlying model sampling only to extract rules [Hai16].

**Explainable AI Visualization** – The visualization of explainable AI methods allows to relate data, model, and output to each other. Ideally, an analyst visually explores and interprets the model training, data, and output to interactively extract rule sets from all components individually. Especially, such visual analysis allows to analyze subsets of the model space as well as the underlying data and helps to identify essential complex hidden relationships (e.g., multilinear dependencies) in the data. A particular visual finding allows an analyst to extract and increase the trust in specific rules. These rules can be verified by visually investigating further trained models and the underlying data. The interpretation and verification of hidden relationships in the data always require domain knowledge. The fully automated extraction and analysis of such rule sets are hardly possible as the training of most of the generated black-box models is non-deterministic. Meaning the rule sets for different models will have a particular overlap, yet not all rules the models learn are reflections of the modeled real-world problem (false positives). For instance, the extraction of linear models using LIME [RSG16] enables to further select and prune rule sets into a reliable subset of rules [RTC18].

*In our exemplary use case*, we got to the XAI step (Fig. 2 (2)). We first train a decision rule set [LBL16] (XAI methods) classifier on the time series (data) to get a basic rule set of the data. Next, we train our k-Nearest baseline on the data and get an F1-Score (metric) of 0.75 on real-world data. Next, we train our chosen

RNN (model) on the anomaly detection task. Our first architecture does not achieve our desired performance on the metric after 40 epochs due to overfitting and a large capacity. After shrinking the architecture, our RNN improves at ten epochs and has a favorable performance after 40 epochs due to better generalizability. We use RxREN [GAK12] (XAI methods) to extract rules from the model at epochs 10 and 40 about the anomaly prediction (output). The decision rule set of the data and extracted RxREN rules of both epochs formulate together our extracted rules for the next step.

### 3.3. Adjusting Extracted Rule Sets

Decision rule sets are the most promising coverage of data and model in general as they formulate a comprehensible and interpretable abstraction for humans [LBL16]. However, the rigorous extraction of rule sets from real-world problems is challenging as the task requires correct as well as complete data, an unbiased and stable model, and reliable extraction of rule sets. Often in real-world environments, the input data does not capture all the required information, the model has some errors, and the extracted rule sets are often hard to obtain. Therefore, a domain expert has to interactively verify as well as adjust the rule sets to confirm the extracted rules with several models. The extracted rule sets are nested and complex if-else statements, which are also machine-readable.

*Our use case* has three kinds of decision sets extracted from the data and the models with specific algorithms mentioned before. Based on these extracted rules, we can start to compare rules from the data and the models. Through iterative and interactively inspecting, pruning, and changing rules, we can verify rules based on our domain knowledge about the data and task (see Fig.2 (3)). We first deep dive into the extracted rules of the second RNN while training to analyze overlapping rules. These rules are the core of the prediction of most of the samples as our RNN has an F1-Score of 0.8 on real-world data after ten epochs. By further inspection of these rules, we find a certain periodicity in one of the attributes to classify one class. Two rules describe such a periodicity by stating two neighboring time ranges and a particular attribute value.

### 3.4. Resulting Model Specification

The model specification is a machine-readable description and describes how the model solves a particular problem, behaves for particular inputs, and can be used to compare the generated knowledge about the model with the initial formulated problem. Generated rule sets contain parts of a model specification as the selection of rules for the specification heavily depends on the underlying data, application domain, and domain expert. A model specification should also be verified collaboratively from multiple persons to enable feedback for refinements and to minimize the bias influenced by single domain experts. The primary purpose of the approach is to describe and examine the solution of the problem by investigating the automatically learned subproblems. The model specification can be used to confirm assumptions about a specific domain problem, compare the generated rules with domain expert knowledge, and to evaluate the usefulness of a model. Overall, the model specification aims to divide complex formulated problems into subproblems to enable domain experts to find transparent and straightforward solutions for the original problem.

*Our use case example* results with a model specification based on the rules our analyst left after the last step. The extracted rules from the model show that the model learned some domain knowledge by detecting certain periodicity and specific frequencies significant for the anomaly. Our extracted model specification now answers our twofold problem formulation (forecasts failures by self-extracted domain knowledge) with:

- F1-Score: 0.75 (baseline) - 0.95 (RNN) on real-world data
- If the time attribute is larger equal to 50, smaller equal to 55, and the att1 attribute is larger equal to 0.95, then there is an anomaly.

The second rule is one of the selected rules of the domain expert to show an example of the domain knowledge.

## 4. Discussion and Opportunities

ModelSpeX is a workflow that characterizes the human-in-the-loop and XAI concepts for the evaluation and verification in the visual analytics process. The workflow specifies XAI methods for the rule set extraction as an important subprocess in the visual analysis of black-box models and facilitates to tackle multiple challenges: the extracting knowledge from non-deterministic methods, evaluation of models, and the integration of a human user in the knowledge generation process. The application of non-deterministic methods (e.g., LIME [RSG16], Anchors [RTC18]) on non-deterministic black-box models poses some difficulties. For example, the completeness of extracted rules, the extraction of false-positive rules, the granularity of the extracted rules, and the evaluation of the quality of the rule sets. In our workflow, the human-in-the-loop has to deal with such challenges and verify extracted rules. A possible solution for such challenges is to guide users in the sense-making process, for instance, making suggestions based on the current extracted rules can help analysts to automate the extraction of rules sets. Another question that arises is, if a list of rules as a model specification is valuable, and thus if the extracted rule sets are able to be mapped directly to gained knowledge. We argue that such specifications are useful as they can help to identify problems with the model, highlight differences to the mental model of the user, and also evaluate the quality of the extracted explanations. However, due to the potential bias of the user, models of different users will likely include different rule sets. We see distinct model specifications as an advantage of our workflow as the comparison and evaluation of different rule sets allows discussions and further analysis of the potential disambiguity of analytical problems.

We want to demonstrate the applicability of the workflow on various real-world example ranging from trajectory forecasting to animal behaviour modeling. Especially, methods for the rule extraction of special applications and interactive visualizations for large decision rule sets are planned to be worked on. These rule extraction algorithms will focus on critical domains like healthcare to introduce a methodology to use complex models in such domains.

### Acknowledgement

This work was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2117 - 422037984 and the European Union's Horizon 2020 research and innovation programme under grant agreement No 830892.

## References

- [AB18] ADADI A., BERRADA M.: Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access* 6 (2018), 52138–52160. doi:10.1109/ACCESS.2018.2870052. 1, 2
- [BBSS03] BEDERSON P., BEDERSON B., SHNEIDERMAN B., SHNEIDERMAN U.: *The Craft of Information Visualization: Readings and Reflections*. Interactive Technologies. Elsevier Science, 2003. 3
- [CHH\*19] CASHMAN D., HUMAYOUN S. R., HEIMERL F., PARK K., DAS S., THOMPSON J., SAKET B., MOSCA A., STASKO J., ENDERT A., GLEICHER M., CHANG R.: A user-based visual analytics workflow for exploratory model analysis. *Computer Graphics Forum* 38, 3 (2019), 185–199. arXiv:1809.10782, doi:10.1111/cgf.13681. 2
- [CS96] CRAVEN M. W., SHAVLIK J. W.: Extracting tree-structured representations of trained networks. *Advances in Neural Information Processing Systems* 8 (1996), 7. 3
- [CSKX15] CHEN C., SEFF A., KORNHAUSER A., XIAO J.: Deepdriving: Learning affordance for direct perception in autonomous driving. In *IEEE International Conference on Computer Vision* (2015), pp. 2722–2730. 1
- [DK17] DOSHI-VELEZ F., KIM B.: Towards A Rigorous Science of Interpretable Machine Learning. *arXiv preprint arXiv:1702.08608* (2017). doi:10.1016/j.bbrc.2004.04.155. 1, 2
- [GAK12] GETHSIYAL AUGASTA M., KATHIRVALAVAKUMAR T.: Reverse engineering the neural networks for rule extraction in classification problems. *Neural Processing Letters* 35, 2 (2012), 131–150. doi:10.1007/s11063-011-9207-8. 3, 4
- [GMR\*18] GUIDOTTI R., MONREALE A., RUGGIERI S., TURINI F., GIANNOTTI F., PEDRESCHI D.: A Survey Of Methods For Explaining Black Box Models. *ACM Computing Surveys* 51, 5 (2018), 93:1—93:42. doi:10.1145/3236009. 1, 2
- [Gol96] GOLEA M.: On the complexity of rule extraction from neural networks and network-querying. In *Rule Extraction From Trained Artificial Neural Networks Workshop, Society For the Study of Artificial Intelligence and Simulation of Behavior Workshop Series (AISB)* (1996). 3
- [Hai16] HAILESILASSIE T.: Rule Extraction Algorithm for Deep Neural Networks: A Review. *International Journal of Computer Science and Information Security* 14, 7 (Sept. 2016), 376–381. 1, 3
- [HMKL18] HOFFMAN R. R., MUELLER S. T., KLEIN G., LITMAN J.: Metrics for Explainable AI: Challenges and Prospects. *arXiv preprint arXiv:1812.04608* (2018). 1
- [KAF\*08] KEIM D., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: Visual analytics: Definition, process, and challenges. In *Information Visualization*. Springer, 2008, pp. 154–175. 2, 3
- [KKEM10] KEIM D., KOHLHAMMER J., ELLIS G., MANSMANN F.: *Mastering the Information Age : Solving Problems with Visual Analytics*. Goslar: Eurographics Association, 2010. 3
- [LBL16] LAKKARAJU H., BACH S. H., LESKOVEC J.: Interpretable decision sets: A joint framework for description and prediction. *International Conference on Knowledge Discovery and Data Mining 13-17-Augu* (2016), 1675–1684. doi:10.1145/2939672.2939874. 3, 4
- [Lip18] LIPTON Z. C.: The myths of model interpretability. *Queue* 16, 3 (Sept. 2018), 30:31—30:57. doi:10.1145/3233231. 2
- [LWLZ17] LIU S., WANG X., LIU M., ZHU J.: Towards Better Analysis of Machine Learning Models: A Visual Analytics Perspective. *Visual Informatics I*, 1 (2017). doi:10.1016/j.visinf.2017.01.006. 1, 2, 3
- [Mob02] MOBLEY R. K.: *An Introduction to Predictive Maintenance*, second ed. ed. Plant Engineering. Butterworth-Heinemann, Burlington, 2002. doi:10.1016/B978-075067531-4/50000-2. 2
- [MZR18] MOHSENI S., ZAREI N., RAGAN E. D.: A Survey of Evaluation Methods and Measures for Interpretable Machine Learning. *arXiv preprint arXiv:1811.11839* (2018). 1
- [Nor88] NORMAN D. A.: *The Psychology of Everyday Things*. Basic books, 1988. 2, 3
- [PF91] PIATESKI G., FRAWLEY W.: *Knowledge Discovery in Databases*. MIT press, 1991. 2, 3
- [RSG16] RIBEIRO M. T., SINGH S., GUESTRIN C.: "Why Should I Trust You?". In *International Conference on Knowledge Discovery and Data Mining* (New York, New York, USA, 2016), ACM Press, pp. 1135–1144. doi:10.1145/2939672.2939778. 1, 3, 4
- [RTC18] RIBEIRO M., TULIO S. S., CARLOS GUESTRIN: Anchors: High-Precision Model-Agnostic Explanations. In *AAAI Conference on Artificial Intelligence* (Oct. 2018). 1, 3, 4
- [Rud18] RUDIN C.: Please Stop Explaining Black Box Models for High Stakes Decisions. *arXiv preprint arXiv:1811.10154* (Nov. 2018). 1
- [SAE\*19] SCHLEGEL U., ARNOUT H., EL-ASSADY M., OELKE D., KEIM D. A.: Towards a Rigorous Evaluation of XAI Methods on Time Series. In *ICCV Workshop on Interpreting and Explaining Visual Artificial Intelligence Models* (2019). 2
- [SHJ\*19] SLACK D., HILGARD S., JIA E., SINGH S., LAKKARAJU H.: How can we fool LIME and SHAP? Adversarial Attacks on Post hoc Explanation Methods. *arXiv preprint arXiv:1911.02508* (2019). 1
- [SSSE19] SPINNER T., SCHLEGEL U., SCHÄFER H., EL-ASSADY M.: explAiner: A Visual Analytics Framework for Interactive and Explainable Machine Learning. *IEEE Transactions on Visualization and Computer Graphics* (2019). 2
- [SSZ\*17] SACHA D., SEDLMIR M., ZHANG L., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing* (2017). doi:10.1016/j.neucom.2017.01.105. 2, 3
- [WSC\*16] WU Y., SCHUSTER M., CHEN Z., LE Q. V., NOROUZI M., MACHEREY W., KRIKUN M., CAO Y., GAO Q., MACHEREY K., KLINGNER J., SHAH A., JOHNSON M., LIU X., ŁUKASZ KAISER, GOUS S., KATO Y., KUDO T., KAZAWA H., STEVENS K., KURIAN G., PATIL N., WANG W., YOUNG C., SMITH J., RIESA J., RUDNICK A., VINYALS O., CORRADO G., HUGHES M., DEAN J.: Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016). URL: <http://arxiv.org/abs/1609.08144>. 1
- [ZMJ16] ZILKE J. R., MENCÍA E. L., JANSSEN F.: DeepRED: Rule extraction from deep neural networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9956 LNAI (2016), 457–473. doi:10.1007/978-3-319-46307-0\_29. 3