

Towards Quantification of Explainability in Explainable Artificial Intelligence Methods

Sheikh Rabiul Islam, William Eberle, Sheikh K. Ghafoor

Department of Computer Science, Tennessee Technological University
Cookeville, Tennessee, United States
sislam42@students.tntech.edu, weberle@tntech.edu, sghafoor@tntech.edu

Abstract

Artificial Intelligence (AI) has become an integral part of domains such as security, finance, healthcare, medicine, and criminal justice. Explaining the decisions of AI systems in human terms is a key challenge—due to the high complexity of the model, as well as the potential implications on human interests, rights, and lives. While Explainable AI is an emerging field of research, there is no consensus on the definition, quantification, and formalization of explainability. In fact, the quantification of explainability is an open challenge. In our previous work (Islam et al. 2019), we incorporated domain knowledge for better explainability, however, we were unable to quantify the extent of explainability. In this work, we (1) briefly analyze the definitions of explainability from the perspective of different disciplines (e.g., psychology, social science), properties of explanation, explanation methods, and human-friendly explanations; and (2) propose and formulate an approach to quantify the extent of explainability. Our experimental result suggests a reasonable and model-agnostic way to quantify explainability.

Introduction

The use of Artificial Intelligence (AI) has spread into a spectrum of high-stakes applications from different domains such as security, finance, healthcare, medicine, and criminal justice, impacting human interests, rights, and lives. Following the pace of the demand, the successful AI-based models have also become so complex that their decisions have become too complicated to express in human terms. This further complicates their adoption in many sensitive disciplines, raising concerns from the ethical, privacy, fairness, and transparency perspectives. The root cause of the problem is the lack of explainability and/or interpretability of the decision.

We have seen the very confusing use of the terms “explainability” and “interpretability” throughout the literature. Some treat these as the same and stick to one, while some differentiate among the two, and others use them ambiguously. Research in Explainable Artificial Intelligence (XAI) is seeing a resurgence after three decades of slowed progress

since some preliminary work on expert systems in the early 1980’s. Currently, there is no consensus on the definition these terms. In addition, the quantification of explainability is another open challenge, which will be difficult until there is a consensus on the concrete definition of those terms (e.g., explainability, interpretability).

In our previous work (Islam et al. 2019) and (Islam et al. 2020), we proposed an approach to incorporate domain knowledge in the model to make the prediction more explainable for two different domain. However, we were unable to quantify the quality of explanations (i.e., validation of explainability). There has been some previous work (Molnar, Casalicchio, and Bischl 2019) that mentions three notions for quantification of explainability. The first two notions involve experimental studies with humans (e.g., domain expert, layperson) that mainly investigate whether a human can predict the outcome of the model (Dhurandhar et al. 2017), (Friedler et al. 2019), (Huysmans et al. 2011), (Poursabzi-Sangdeh et al. 2018). However, the third notion (proxy tasks) does not involve a human. Instead they use known truth as a metric (e.g., the less the depth of the decision tree, the more explainable the model).

In this work, our contributions are as follows:

1. We define explainability and interpretability from a multi-domain perspective and clarify their resemblance.
2. We analyze the properties of explanation, explanation methods, and human friendly explanations.
3. We present a potential way to formalize and quantify explainability with experimental results using proxy tasks.

Explainability, and Interpretability

(Miller 2018) argue that most of the work on XAI focuses on the researcher’s intuition of what constitutes a good explanation. However, there exists a vast area of research in philosophy, psychology, and cognitive science on how people generate, select, evaluate, and represent explanations and associated cognitive biases and social expectations towards the explanation process. Therefore, the author emphasizes that the research on Explainable AI should incorporate studies from these different domains.

From the social science perspective, according to (Miller 2018) and (Lombrozo 2006), explanation is both a process and product:

- Cognitive explanation process: identifies the causes for an event, perhaps concerning particular counterfactual cases, and a subset of these causes is selected as the explanation.
- Product: the explanation that results from the cognitive explanation process.
- Social process: the process of transferring knowledge between explainer and explainee, generally an interaction between groups of people, in which the goal is to provide the explainee with enough information to understand the causes of the event.

This definition of explanation insists that it is both a process and product (i.e., outcome), and focuses on understanding the causes of the event.

Furthermore, according to psychologist (Lombrozo 2006), “Explanations are ... the currency in which we exchanged belief”. This definition stresses the need for high fidelity for the explanation, in other words, gaining trust from the explanation recipient. Moreover, according to physicist (Deutsch 1998) who is the pioneer of quantum computation: *explanations* consist of *interpretations* of how the world works and why. This definition suggests that explanations are roughly equal to or a superset of interpretations. Interpretations are the building blocks of explanations; without interpretation, explanation is incomplete. Furthermore, the terms “explainability” and “interpretability” are simply the extent of explanation and interpretation accordingly.

There are other mentionable definitions too. According to (Miller 2018), “interpretability is the degree to which a human can understand the cause of a decision”. Besides, according to (Kim et al. 2017) “Interpretability is the degree to which a human can consistently predict the model’s result”. Furthermore, according to (Ribeiro, Singh, and Guestrin 2016) “By ‘explaining a prediction’, we mean presenting textual or visual artifacts that provide qualitative understanding of the relationship between the instance’s components (e.g. words in text, patches in an image) and the model’s prediction”. In other words, (Ribeiro, Singh, and Guestrin 2016) put an emphasis on the qualitative understanding of the relationship between input and output (i.e., a selective/suitable set of information pieces that together can refer to a cause to an event, in contrast to the complete causal attributions).

To this end, in the case of an intelligent system (i.e., AI/ML-based system), it is evident that explainability is more than interpretability in terms of importance, completeness, and fidelity of prediction. Based on that, we choose to keep our focus entirely on explainability instead of interpretability. Finally, analyzing different definitions from the literature, we come up with following:

Definition: *Explainability* of an AI model’s prediction is the extent of transferable *qualitative understanding* of the relationship between model input and prediction (i.e., selective/suitable causes of the event) in a recipient friendly manner.

We proceed further based on this informal definition of explainability.

Background

There are two primary directions of research towards evaluation of explainability of an AI/ML model: (1) model complexity-based and (2) human study-based.

Model Complexity-based Explainability Evaluation

In the literature, model complexity and (lack of) model interpretability are often treated as the same (Molnar, Casalicchio, and Bischl 2019). For instance, in (Fürnkranz, Gamberger, and Lavrač 2012) and (Yang, Rudin, and Seltzer 2017), model size is often used as a measure of interpretability (e.g., number of decision rules, depth of tree, number of non-zero coefficients).

(Yang, Rudin, and Seltzer 2017) propose a scalable Bayesian Rule List (i.e., probabilistic rule list) consisting of a sequence of IF-THEN rules, identical to a decision list or one-sided decision tree. Unlike the decision tree that uses greedy splitting and pruning, their approach produces a highly sparse and accurate rule list with a balance between interpretability, accuracy, and computation speed.

Furthermore, according to (Rüping and others 2006), while the number of features and the size of the decision tree are directly related to the interpretability, the optimization of the tree size or features (i.e., feature selection) is costly as it requires generation of a large set of models and their elimination in subsequent steps. Besides, reducing the tree size (i.e., reducing complexity) increases error; however, they could not find a way to formulate the relation in simple functional form.

More recently, (Molnar, Casalicchio, and Bischl 2019) attempt to quantify the complexity of the arbitrary machine learning model with a model agnostic measure. In that work, the author demonstrates that when the feature interaction (i.e., the correlation among features) increases, the quality of representations of explainability tools degrades. For instance, the explainability tool ALE Plot starts to show harsh lines (i.e., zigzag lines) as feature interaction increases. In other words, with more interaction comes a more combined influence in the prediction, induced from different correlated subsets of features (at least two), which ultimately makes it hard to understand the causal relationship between input and output, compared to an individual feature influence in the prediction. In fact, from our study of different explainability tools (e.g., LIME, SHAP, PDP), we have found that the correlation among features is a key stumbling block to represent feature contribution in a model agnostic way. Keeping the issue of feature interactions in mind, (Molnar, Casalicchio, and Bischl 2019) proposes a technique that uses three measures: number of features, interaction strength among features, and the main effect (excluding the interaction part) of features to measure the complexity of a post-hoc model for interpretation. Although, their work mainly focuses on model complexity for post-hoc models, that acted as a precursor to formulate our approach of explainability quantification. Our approach to quantify explainability is model ag-

nostic and is for a model of any notion (e.g., pre-modeling, post-hoc).

Human Study-based Explainability Evaluation

The following work deals with the application and human-level evaluation of explainability involving human studies.

(Huysmans et al. 2011) investigate the suitability of different alternative representation formats (e.g., decision tables, (binary) decision trees, propositional rules, and oblique rules) for classification tasks primarily focusing on the explainability of results rather than accuracy or precision. They discover that decision tables are the best in terms of accuracy, response time, confidence of answer, and ease of use.

Furthermore, (Dhurandhar et al. 2017) argue that interpretability is not an absolute concept; rather, it is relative to the target model, and may or may not be relative to the human. Their finding suggests that a model is readily interpretable to a human when it uses no more than seven pieces of information (Miller 1956). Although, this might vary in task to task and person to person. For instance, a domain expert might consume a lot more detail information depending on their experience.

The work of (Poursabzi-Sangdeh et al. 2018) is a human-centered approach, focusing on previous work on human trust in a model from psychology, social science, machine learning, and human-computer interaction communities. In their experiment with human subjects, they vary factors (e.g., number of features, whether the model internals is clear or a black box) that makes a model more or less interpretable and measures how the variation impacts the prediction of human subjects. Their result suggests that participants who were shown a clear model with a small number of features were more successful in simulating the model's predictions and trusted the model's predictions.

Furthermore, (Friedler et al. 2019) investigate interpretability of a model based on two of its definitions: simulatability, which is a user's ability to predict the output of a model on a given input; and "what if" local explainability, which is a user's ability to predict changes in prediction in response to changes in input, given the user has the knowledge of a model's original prediction for the original input. They introduce a simple metric called runtime operation count that measures the interpretability, that is, the number of operations (e.g., the arithmetic operation for regression, the boolean operation for trees) needed in a user's mind to interpret something. Their findings suggest that the interpretability decreases with an increase in the number of operations.

Properties of Explanations and Explanation Methods

Although our main interest lies in explainable prediction, it depends on the model that generates the explanations. So, it is crucial to analyze the properties of both the explanation and the explanation method that generates the explanations. (Robnik-Šikonja and Bohanec 2018) and (Molnar and others 2018) attempt to define the properties of explanation

methods and individual explanations, which we present in the following sections.

Properties of Explanation Methods

Some properties of *explanation methods* include:

- **Expressive Power:** Refers to the structure (e.g., decision trees, IF-THEN rules, weighted sum, and natural language) of the explanation method.
- **Translucency:** The level in which the explanation method relies on looking into (e.g., coefficient for a linear regression, node splitting point in the tree-based approach) the ML model. For instance, explanation methods that rely on intrinsically interpretable methods like linear regression are highly translucent; however, counterfactual model agnostic explanation methods that leverage the changes in output in response to input for explanation have zero translucency as it does not look at model at all.
- **Portability:** The number of ML models that the explanation method covers. Usually low translucency (i.e., ML model usually remains black box) comes with more portability.
- **Algorithmic Complexity:** The required time to generate the explanation.

Properties of Individual Explanations

Some properties of *individual explanations* include:

- **Accuracy:** The explanation needs to be accurate enough when fidelity is essential. Although, low accuracy of explanation might be a cause of the low accuracy of the ML/AI model.
- **Fidelity:** How well the explanation approximates the prediction of the black-box model. Usually a highly accurate model with high fidelity leads to highly accurate explanations. Thus, fidelity and accuracy are related.
- **Consistency:** Refers to the extent of consistency among explanations for different ML models on the same task.
- **Stability:** Refers to the extent of similarity of explanations for similar instances. While consistency compares explanations between different models, stability compares the explanations of similar instances for a particular model.
- **Comprehensibility:** The extent to which the recipient of explanation understands the explanation, which is very hard to measure. A few measures could be the number of features with non-zero weights in a linear model, or the number of rules in a tree. Usually a human can comprehend 7+-2 pieces of information at a time (Miller 1956)
- **Certainty:** Many ML models provide the probability of the target class. Similarly, explanations with a certainty value are expected to be useful.
- **Degree of Importance:** How well the explanation covers the important features or how well the explanation reflects parts of the explanations. For example, from the decision rules, we can understand which of the rules are more important (e.g., important features are at the beginning of a ruleset, or the top of the tree).

Recent laws (Goodman and Flaxman 2016) (Wyden) focus on the impact of an algorithmic decision on human rights, interests, and lives. In that sense, we need a human-friendly explanation, which has slightly different properties than the explanation that does not consider humans as a key factor.

Properties of Human-friendly Explanation

According to (Miller 2018), humans usually prefer short explanations that contrast the current situation with a situation in which that event would not have occurred (a.k.a., counterfactual explanations). Furthermore, the human friendly explanation does not consider all factors (i.e., selective in nature) for a particular prediction or behavior. However, if one needs to legally specify all influencing factors or need to debug the machine learning model, there is a need for a complete causal attribution (Molnar and others 2018), which is out of the scope of human-friendly explanation. That is the reason behind mentioning the term *qualitative* instead of *quantitative* in our definition of explanation (see Definition of Explainability).

In summary, according to (Štrumbelj and Kononenko 2011), (Miller 2018) (Lipton 1990), and (Molnar and others 2018), human friendly explanations are: *contrastive*—why a prediction (e.g., loan was rejected) was made instead of the alternative prediction (e.g., loan was accepted), *selected*—does not cover the complete list of causes of an event, *social*—social context (e.g., explaining to a layperson or domain expert) determines the nature of the explanations, *abnormal behavior focused*—when a criteria (e.g., a particular feature value) is rare and has influence in the prediction, then it should be included (i.e., should have higher precedence in case of other criteria with the same influence) in the explanation, *truthful*—explanation should be as truthful as possible, although selectiveness comes first which might exclude some of the true reasons, *consistent*—consistent with prior belief, and *general and probable*—good explanations are general and probable, although this contradicts with the claims that abnormal causes make good explanations, abnormal causes have higher preference over general and probable explanations (see *abnormal behavior focused* point).

Explainability Quantification Method

Usually, humans can relate and process 7+2 pieces of information (i.e., cognitive chunks) to understand something (Miller 1956). For instance, suppose that, in the most generalized form, the quality of an explanation is dependent upon the number of cognitive chunks that the recipient has to relate to in order to understand an explanation (i.e., the less, the better). Lets assume, E = explainability; N_c = number of cognitive chunks; I = interaction; N_i = number of input cognitive chunks; and N_o = number cognitive chunks involved in the explanation representation (i.e., output cognitive chunks).

$$E = \frac{1}{N_c} \quad (1)$$

However, sometimes, these cognitive chunks are correlated and their influence/contribution/abilities are not mutually

exclusive. This interaction among cognitive chunks complicates the explainability. So we penalize Formula 1 for having an interaction among cognitive chunks, resulting in Formula 2.

$$E = \frac{1}{N_c} + (1 - I) \quad (2)$$

Where, the interaction I ranges in between 0 and 1, and the less the interaction, the better the explainability, so we take the complement of that.

Formula 2 is in a form that can be applied to any of the application, domain, or proxy level explanation evaluations described before. However, from the perspective of an application and domain level evaluation of explainability, to progress further, we need human studies that are out of the scope of this work. Instead, in this work, we focus on the proxy level evaluation of explainability that considers different properties of output representation (e.g., depth of decision tree, length of rule list) as a metric for evaluation.

Furthermore, we need to breakdown the number of cognitive chunks more to get a better evaluation. Both the number of input cognitive chunks in the model and the number of output cognitive chunks involved in the representation of output are important to understand the causal relationship, which is vital for explanation. While the ideal explainability case would be when there is only one input and one output cognitive chunk (no chance of interaction), that is unusual in real-world situations. Following the segregation of input and output cognitive chunks, Formula 2 can be re-written as Formula 3:

$$E = \frac{1}{N_i} + \frac{1}{N_o} + (1 - I) \quad (3)$$

where N_i refers to the number of input cognitive chunks and N_o refers to the number cognitive chunks involved in the explanation representation (i.e., output cognitive chunks). Usually, the more these cognitive chunks, the more complicated the explanation becomes. So, the ratio of the best case (i.e., one cognitive chunk) and observed case is added towards total explainability.

Also, Formula 3 has three predicates, which might have different influences on the quantification of explainability, and different domains might have different implications (e.g., accuracy vs explainability trade off). So, we add a weight term with each of the predicates, considering the weights are constant (e.g., .3333) by default, and their summation is equal to 1: $w_1 + w_2 + w_3 = 1$. However, these weights can be set to a different distribution, perhaps dependent upon a particular domain (e.g., healthcare, finance).

After the addition of the weight terms, Formula 3 becomes Formula 4:

$$E = \frac{w_1}{N_i} + \frac{w_2}{N_o} + w_3(1 - I) \quad (4)$$

Formula 4 can then be used to quantify the explainability of the explanation method (i.e., global explainability). We can use Formula 4 to also quantify the explainability of an instance level prediction (i.e., local explainability). In that case, the first predicate of Formula 4 (including the

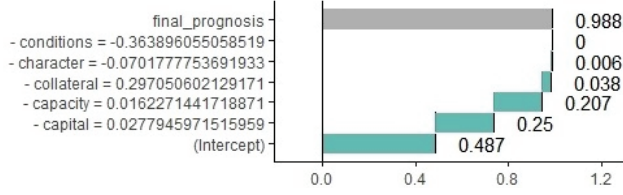


Figure 1: Breakdown of feature contributions for a random sample.

weight term) remains the same (i.e., the same number of input chunks). However, predicate 2 and predicate 3 will be different from instance to instance as a different set of cognitive chunks with different interaction strengths might be involved in the representation of explanation for a particular instance as explanations are selective.

Experimental Results

We provide a brief overview of the experimental settings of our previous work (Islam et al. 2019) that we will then use for our proposed explainability quantification method. In this example, we incorporate domain knowledge in the model for mortgage bankruptcy prediction, and apply different supervised ML algorithms in three different ways:

1. Using *original features*: we use all original features that have a non-zero effect on prediction (features with zero effect were removed in a data pre-processing stage by observing the Random Forest’s feature importance).
2. Using *domain-related features*: we select a subset of features that match with extracted domain knowledge (e.g., 5 C’s of credit) from the domain. Traditionally, credit risk is assessed using the 5C’s of credit (character, capital, capacity, collateral, and cycle), a popular domain principle to for determining credit risk. We extract necessary domain knowledge and map to the 5 C’s of credit to get the domain features (Islam et al. 2019).
3. Using *newly constructed features*: we construct a very concise set of new features (i.e., five features—one feature for each element of the domain principle 5C’s of credit) from the domain-related features using the quantitative measure of gain/compromises (i.e., the cumulative sum of related feature values times the correlation coefficient) associated with each element of the domain principle.

Furthermore, we represent the predicted output as a composition of individual elements of the domain principle 5C’s of credit (i.e., *newly constructed features*) (Figure 1).

One will notice that the representation (Figure 1) of the prediction in terms of the *newly constructed features* provides better explainability as the final prediction is segregated into the individual influences of a very concise set of intuitive features (i.e., five compared to 30). However, there is no way to quantify the level of perceived explainability. To use our proposed formula in this work (Formula 4) to quantify explainability, we need to calculate the *interaction strength* (I) too. We measure the interaction strength among

	Original	Domain	Constructed
Input chunks (N_i)	30	7	7
Output chunks (N_o)	30	7	5
Int. Strength (I)	0.556	0.5233	0.5251
Explainability (E)	0.1701	0.2539	0.2723

Table 1: Comparison of explainability

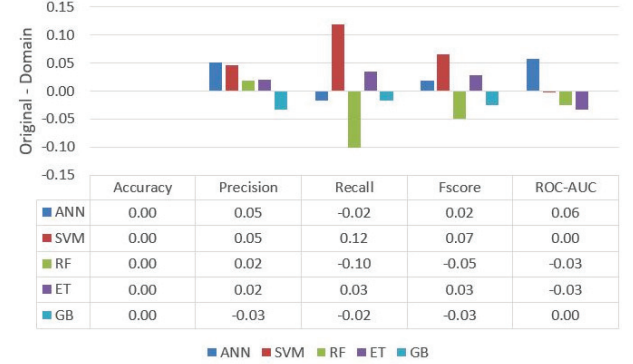


Figure 2: Dispersion in performance—original features minus domain-related features

features using R’s iml package that uses the partial dependence of an individual features as the basis for calculating *interaction strength* (I).

Applying Formula 4, on metadata (Table 1) of three different feature settings, we see that *newly constructed features* (5 C’s of credit) provide the best explainability score of .2723, which is an improvement of 60.14% compared to the 0.17 that we get using the *original features* (Table 1). In fact, even if we apply the state-of-the-art methods of post-hoc interpretability/explainability like SHAP, the explainability will be still limited to 0.17 as it does not reduce the number of cognitive chunks to represent output. Besides, using *domain related features*, the explainability score is .2538, which is better than using the original features, although a little worse than using the *newly constructed features*. We also applied Principal Component Analysis (PCA) to see the extent of compromise in information in exchange for achieved explainability in domain mapped and newly constructed features. We found that PCA takes the first 29 principal components (i.e., new features) to get a similar performance to when we use the original features (30 cognitive chunks). Therefore we can say that the compromise in information for explainability is very negligible. Besides, the principal components of PCA lacks explainability.

Furthermore, this explainability gain comes with a negligible cost of performance (Figure 2, Figure 3) which is expected and a known trade off for explainability methods. In fact, for a few algorithms (e.g., Random Forest (RF), Gradient Boosting (GB)), the *newly constructed features*, and *domain related features* lead to better *recall*, which is crucial for anomaly detection where the target class instances are very few compared to non-target class instances. In that

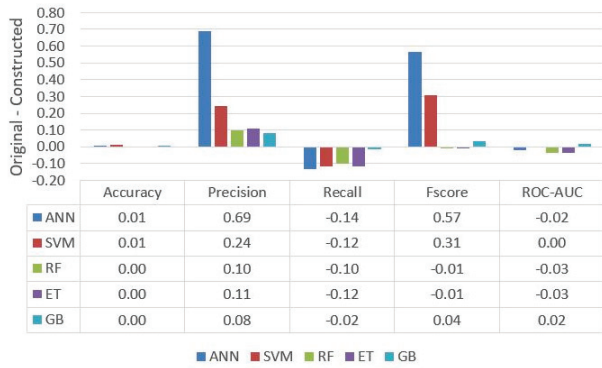


Figure 3: Dispersion in performance—original features minus newly constructed features

sense, besides a better explainability, the domain knowledge also benefits us in terms of improved performance (e.g., better recall, less features leads to less computation time).

Conclusion

Explainable decisions from commercial AI systems are going to be a standard imposed by regulators to eliminate bias and discrimination, and ensure trust. Our work attempts to establish a concrete definition of explainability, properties of explanations and explanation methods, and a way to quantify explainability which is transferable to a variety of explainability methods or tools. For the quantification of explainability, we only present an approach for the proxy method. As an extension of this work, we would like to address human studies, and investigate the effectiveness among different approaches (e.g., supervised, unsupervised, semi-supervised) for different application areas (e.g., natural language processing, image recognition).

Acknowledgment

Thanks to Tennessee Tech’s College of Engineering (CoE) and Cyber-security Education, Research and Outreach Center (CEROC) for supporting this research.

References

Deutsch, D. 1998. *The fabric of reality*. Penguin UK.

Dhurandhar, A.; Iyengar, V.; Luss, R.; and Shanmugam, K. 2017. Tip: Typifying the interpretability of procedures. *arXiv preprint arXiv:1706.02952*.

Friedler, S. A.; Roy, C. D.; Scheidegger, C.; and Slack, D. 2019. Assessing the local interpretability of machine learning models. *arXiv preprint arXiv:1902.03501*.

Fürnkranz, J.; Gamberger, D.; and Lavrač, N. 2012. Rule learning in a nutshell. In *Foundations of Rule Learning*. Springer. 19–55.

Goodman, B., and Flaxman, S. 2016. Eu regulations on algorithmic decision-making and a “right to explanation”. In *ICML workshop on human interpretability in machine learning (WHI 2016)*, New York, NY.

Huysmans, J.; Dejaeger, K.; Mues, C.; Vanthienen, J.; and Baesens, B. 2011. An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models. *Decision Support Systems* 51(1):141–154.

Islam, S. R.; Eberle, W.; Bundy, S.; and Ghafoor, S. K. 2019. Infusing domain knowledge in ai-based “black box” models for better explainability with application in bankruptcy prediction. *arXiv preprint arXiv:1905.11474*.

Islam, S. R.; Eberle, W.; Ghafoor, S. K.; Siraj, A.; and Rogers, M. 2020. Domain knowledge aided explainable artificial intelligence for intrusion detection and response. In *AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering*.

Kim, B.; Wattenberg, M.; Gilmer, J.; Cai, C.; Wexler, J.; Viegas, F.; and Sayres, R. 2017. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). *arXiv preprint arXiv:1711.11279*.

Lipton, P. 1990. Contrastive explanation. *Royal Institute of Philosophy Supplements* 27:247–266.

Lombrozo, T. 2006. The structure and function of explanations. *Trends in cognitive sciences* 10(10):464–470.

Miller, G. A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* 63(2):81.

Miller, T. 2018. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*.

Molnar, C., et al. 2018. Interpretable machine learning: A guide for making black box models explainable. *E-book at https://christophm.github.io/interpretable-ml-book/* 10.

Molnar, C.; Casalicchio, G.; and Bischl, B. 2019. Quantifying interpretability of arbitrary machine learning models through functional decomposition. *arXiv preprint arXiv:1904.03867*.

Poursabzi-Sangdeh, F.; Goldstein, D. G.; Hofman, J. M.; Vaughan, J. W.; and Wallach, H. 2018. Manipulating and measuring model interpretability. *arXiv preprint arXiv:1802.07810*.

Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. Why should i trust you?: Explaining the predictions of any classifier. In *22nd ACM SIGKDD*, 1135–1144. ACM.

Robnik-Šikonja, M., and Bohanec, M. 2018. Perturbation-based explanations of prediction models. In *Human and Machine Learning*. Springer. 159–175.

Rüping, S., et al. 2006. Learning interpretable models.

Štrumbelj, E., and Kononenko, I. 2011. A general method for visualizing and explaining black-box regression models. In *International Conference on Adaptive and Natural Computing Algorithms*, 21–30. Springer.

Wyden, B. Algorithmic accountability.

Yang, H.; Rudin, C.; and Seltzer, M. 2017. Scalable bayesian rule lists. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3921–3930. JMLR. org.

References

- Deutsch, D. 1998. *The fabric of reality*. Penguin UK.
- Dhurandhar, A.; Iyengar, V.; Luss, R.; and Shanmugam, K. 2017. Tip: Typifying the interpretability of procedures. *arXiv preprint arXiv:1706.02952*.
- Friedler, S. A.; Roy, C. D.; Scheidegger, C.; and Slack, D. 2019. Assessing the local interpretability of machine learning models. *arXiv preprint arXiv:1902.03501*.
- Fürnkranz, J.; Gamberger, D.; and Lavrač, N. 2012. Rule learning in a nutshell. In *Foundations of Rule Learning*. Springer. 19–55.
- Goodman, B., and Flaxman, S. 2016. Eu regulations on algorithmic decision-making and a “right to explanation”. In *ICML workshop on human interpretability in machine learning (WHI 2016)*, New York, NY.
- Huysmans, J.; Dejaeger, K.; Mues, C.; Vanthienen, J.; and Baesens, B. 2011. An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models. *Decision Support Systems* 51(1):141–154.
- Islam, S. R.; Eberle, W.; Bundy, S.; and Ghafoor, S. K. 2019. Infusing domain knowledge in ai-based “black box” models for better explainability with application in bankruptcy prediction. *arXiv preprint arXiv:1905.11474*.
- Islam, S. R.; Eberle, W.; Ghafoor, S. K.; Siraj, A.; and Rogers, M. 2020. Domain knowledge aided explainable artificial intelligence for intrusion detection and response. In *AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering*.
- Kim, B.; Wattenberg, M.; Gilmer, J.; Cai, C.; Wexler, J.; Viegas, F.; and Sayres, R. 2017. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). *arXiv preprint arXiv:1711.11279*.
- Lipton, P. 1990. Contrastive explanation. *Royal Institute of Philosophy Supplements* 27:247–266.
- Lombrozo, T. 2006. The structure and function of explanations. *Trends in cognitive sciences* 10(10):464–470.
- Miller, G. A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* 63(2):81.
- Miller, T. 2018. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*.
- Molnar, C., et al. 2018. Interpretable machine learning: A guide for making black box models explainable. *E-book at <https://christophm.github.io/interpretable-ml-book/>*; 10.
- Molnar, C.; Casalicchio, G.; and Bischl, B. 2019. Quantifying interpretability of arbitrary machine learning models through functional decomposition. *arXiv preprint arXiv:1904.03867*.
- Poursabzi-Sangdeh, F.; Goldstein, D. G.; Hofman, J. M.; Vaughan, J. W.; and Wallach, H. 2018. Manipulating and measuring model interpretability. *arXiv preprint arXiv:1802.07810*.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. Why should i trust you?: Explaining the predictions of any classifier. In *22nd ACM SIGKDD*, 1135–1144. ACM.
- Robnik-Šikonja, M., and Bohanec, M. 2018. Perturbation-based explanations of prediction models. In *Human and Machine Learning*. Springer. 159–175.
- Rüping, S., et al. 2006. Learning interpretable models.
- Štrumbelj, E., and Kononenko, I. 2011. A general method for visualizing and explaining black-box regression models. In *International Conference on Adaptive and Natural Computing Algorithms*, 21–30. Springer.
- Wyden, B. Algorithmic accountability.
- Yang, H.; Rudin, C.; and Seltzer, M. 2017. Scalable bayesian rule lists. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3921–3930. JMLR. org.