
A Quantitative Human-Grounded Evaluation Process for Explainable ML

Katharina Beckh*
Fraunhofer IAIS
katharina.beckh@iais.fraunhofer.de

Sebastian Müller*
University of Bonn
semueller@uni-bonn.de

Stefan Rüping
Fraunhofer IAIS
stefan.rueping@iais.fraunhofer.de

*equal contribution

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Copyright held by the owner/author(s).
CHI'22, April 30 – May 6, 2022, New Orleans, USA
ACM 978-1-4503-6819-3/20/04.
<https://doi.org/10.1145/3334480.XXXXXX>

Abstract

Methods from explainable machine learning are increasingly applied. However, evaluation of these methods is often anecdotal and not systematic. Prior work has identified properties of explanation quality and we argue that evaluation should be based on them. In this work, we provide an evaluation process that follows the idea of property testing. The process acknowledges the central role of the human, yet argues for a quantitative approach for the evaluation. We find that properties can be divided into two groups, one to ensure trustworthiness, the other to assess comprehensibility. Options for quantitative property tests are discussed. Future research should focus on the standardization of testing procedures.

Introduction

The development and adoption of complex machine learning (ML) models has given rise to research that seeks to make these models explainable [7, 4, 1]. Explainability methods aim to provide reliable insights into reasons for model behavior. The motivations for why these insights are sought after are manifold: Debugging models during development, finding undesired artifacts in data, learning from models to gain novel scientific insights and to verify that the model fulfills ethical or legal requirements [17]. These insights enable informed decisions about trusting ML models, thus giving a foundation to foster societal acceptance. De-

User Independent Properties

Faithfulness refers to how accurate the explanation is to the model's reasoning.

Completeness describes how much of the model behavior is described.

Consistency refers to how deterministic and implementation-invariant the explanation method is.

Continuity similar inputs should have similar explanations.

Contrastivity answers "why not?" or "what if?" questions.

Confidence probability information in the explanation.

User Dependent Properties

Compactness represents the size of the explanation.

Usability degree of fulfillment of user satisfaction and effectiveness.

Plausibility refers to how convincing an explanation is to a user.

spite the multitude of explanation methods available, one open question is still what evaluation process is suitable to quantify performance of an explanation system [18]. For explainable machine learning to evolve as a field we need to find ways to systematically evaluate explanation methods, taking into account both the user and that explanations represent model behavior [2, 11, 23, 19]. In this work, we present an evaluation process that factors in both goals and describe directions for quantitative testing to facilitate adoption by practitioners.

Background

Evaluation of explainable systems is typically categorized into three types: (1) application-grounded evaluation with humans and applied tasks (2) human-grounded evaluation with humans but proxy tasks and (3) functionally-grounded evaluation without humans and proxy tasks instead [9].

In a recent review of more than 300 papers from prominent ML conferences, it was found that evaluation is roughly 30% anecdotal evidence and roughly 20% with users [18]. In a smaller-scale survey on explainable natural language processing, 50 papers were reviewed [7]. Roughly 60% used no or "informal" evaluation, 23% comparison to ground truth and 17% human evaluation. The percentage of human evaluation is comparable to the more recent study, however, the number for papers with anecdotal evaluation is double. An underlying reason could be that "no or informal" evaluation includes approaches that were deemed as not anecdotal by the other study. These reviews reveal that at least 30% and up to 60% of papers do not present adequate evaluation for explanation quality. Although explanation methods are abundantly available and incorporated in ML systems, the field is still lacking standards regarding the evaluation of explanation quality similar to the standards we expect from an evaluation on predictive performance.

Explanation properties

There exist multiple works describing different properties of explanations that can be used for evaluation [11, 15, 25, 24]. Desired properties and definitions are exemplified in the list in the margin. The list is not exhaustive and what properties (sometimes referred to as attributes or notions [25]) are necessary to consider, depends on the particular application. A frequently discussed distinction is one between evaluating **faithfulness** and **plausibility** of an explanation [11]. Faithfulness is also referred to as correctness or fidelity; plausibility as coherence [18]. Evaluation strategies for these properties differ. For faithfulness, quantitative evaluation metrics exist, e.g. "sensitivity" and "infidelity" [3, 26] for feature attribution methods. Plausibility can be evaluated by comparing ML explanations against "ground-truth" explanations [16]. For an overview of properties we refer to a recent survey which lists evaluation methods for each property [18]. With a number of identified properties, the question arises why we do not put these properties to use to systematically evaluate explainability methods.

Trustworthiness and Comprehensibility

None of the goals of explainable ML, be it generating scientific insight or enabling model verification, can be achieved on the basis of faulty explanations. Hence, it is crucial to ensure that explainability systems deliver a high quality output that represents the model behavior truthfully [2, 11]. At the same time, all explanations are meant to be understood by a human user. Therefore, explanations are only useful if they are comprehensible [23, 19]. An evaluation process for explanation systems will need to ensure both, the quality of the explanations as well as the comprehensibility for the target user. We distinguish between user independent and user dependent properties with respect to the evaluation process (see subsequent sections). For this, we define the notion of **trustworthiness** to encompass all properties that

measure the degree to which the explanations reflect the true model behavior. Analogously, we define **comprehensibility** to encompass all properties that measure to what degree and how efficiently the user is able to reach their interaction goals.

Ensuring Trustworthiness:

User Independent Properties (UIPs)

Human preference for an explanation does not necessarily correlate with its trustworthiness [16]. To protect against user bias, we argue that a positive assessment of UIPs needs to be a pre-requisite for the subsequent user-dependent evaluation of the system.

Not all properties are applicable to all explanation methods. For example, explanation methods that involve random sampling are naturally less consistent than deterministic methods. One idea is to use fact sheets to track applicable properties of an explanation method [24], thus making it easier to gauge what properties are relevant for a given system.

We have the impression that existing evaluation measures for UIPs are not widely used yet, but hope that the integration of measures into emerging coding-libraries¹ has an effect on practice.

Ensuring Comprehensibility:

User Dependent Properties (UDPs)

Carrying out user studies is a costly and time-consuming process. We argue that generating reports on UDPs should be made as easy as possible for practitioners. We thus want to discuss possibilities to facilitate quantitative measurements of UDPs.

¹e.g. [Captum](#) and [Quantus](#)

Compactness is a property that is closely tied to managing users' *Cognitive Load* [20]. For this, we can turn to the concept of *Working Memory* [6] that describes the limit of how many *chunks* (meaningful items/ thoughts) a human can actively entertain at the same time. What constitutes a chunk in the respective application needs context-level definition. For text data it might be a single word, for image data it might be all pixels belonging to the same object. What constitutes a chunk also depends on user experience in the respective domain, e.g. a more experienced user having developed a more complex mental model to associate information with can process larger amounts of information than inexperienced users.

Plausibility measures to what extent the explanation corresponds to user expectations. Providing practitioners with a ground truth dataset containing human verified explanations would enable them to report a degree of explanatory overlap as explored by [16]. Remarkably, the study also found user biases in the subjective rating of explanations. This indicates that purely qualitative evaluation is not sufficient providing another reason to develop quantitative measures. A possible disadvantage of ground truth datasets is that they are task specific, hence, transferability might be an issue. Another direction for plausibility checks are test scenarios, inspired by software engineering practices [22]. Test scenarios can be generated according to user expectations. In a scenario with text data and feature highlighting, possible checks could be the following: "words with low information value, e.g. *the*, should not be part of an explanation" or "replacing adjectives with synonyms should yield the same attribution". However, usefulness of ground truth datasets or other test scenarios have their limits. If an explanation does not comply with user expectations that does not mean that it is not a faithful description of model behavior. Discrepancies between expected explanation and

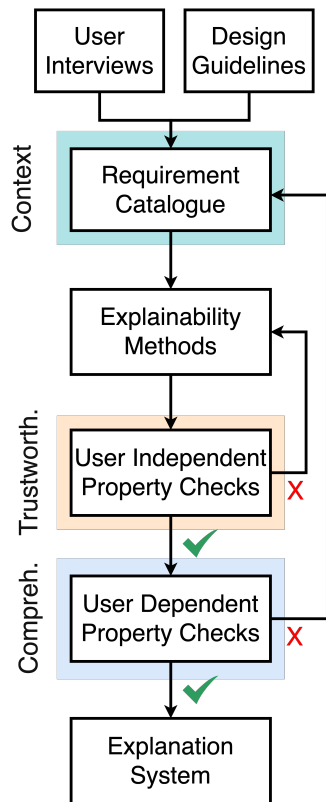


Figure 1: Human-grounded evaluation process for explanation systems.

observed explanations have to be investigated carefully as to not blame the explainability method when it is really the model that behaves unexpectedly. If the explanation is trustworthy, a plausibility check does not reveal anything about the explanation but only something about the model.

Usability is the most elusive property of the three because it not only depends on the explanation methods used but also on presentation, UI, tech-affinity of the user and the specific questions the user wants to be answered. Options to quantitatively measure usability are time-on-task or task-success with an explanation condition (with or without explanation) [12, 8, 21]. The challenge remains that user studies, whether they are qualitative or quantitative, require time and resources.

This list of properties is non-exhaustive and relevance of individual metrics may vary depending on the application setting or the diversity of the targeted user group.

Process for Human-Grounded Quantitative Evaluation

To establish good evaluation practices we propose the evaluation process shown in Figure 1. The key idea of this process is to enforce UIP evaluation both before evaluating UDPs and also after any explanation changing alterations to the system.

At first the **Context** needs to be captured [10]. **User Interviews** provide an investigation into what goals the users want to accomplish and what questions need to be answered, e.g. [13]. Beyond immediate user intent, engineers may rely on established **Design Guidelines** for comparable scenarios to inform initial design decisions, e.g. regarding compositionality.

This informs the **Requirement Catalogue** used to determine what **Explainability Methods** should be selected to fulfill identified user needs. Upon implementation, the system is evaluated with regards to applicable UIPs. If UIPs indicate an improper representation of model behavior by the explanations, the engineer returns to choosing different Explainability Methods and UIP checks are repeated. Once UIP checks are passed, the system is assessed with regards to the UDPs. This can involve a real user study or a quantitative evaluation. If the explanations are deemed incomprehensible, this requires an update of the Requirement Catalogue. Thus, UDPs can impact UIPs: For example, if tests failed for Compactness and the user cannot cope with the amount of information presented, displaying less information has an impact on Completeness of the displayed explanation. If the system is adapted, a re-evaluation of the UIPs is necessary to ensure trustworthiness is maintained.

Most likely, UIP-UDP-trade-offs will be unavoidable in practice, e.g. due to specific user needs or a lack of choice of better explanation methods. Finding a balance will be difficult. Returning to our previous example: A UDP check indicates that the user prefers a more compact explanation. The engineer implements a slider to adjust the amount of chunks to be displayed at the same time. Seemingly, all information is still available to the user, yet the user can now pick at what level of compactness the explanation appears the most plausible. This rather enables them to confirm their own bias of the model instead of investigating the true model behavior. Providing the user with information about how their actions affect trustworthiness might be a possible remedy. However, in some scenarios, e.g. where an explanation system is used to investigate a model for certification or in a legal dispute, trustworthiness demands may be non-negotiable.

In general, we see the risk that compromising UIPs, such as faithfulness, opens the way for manipulation in adversarial settings [5]. We expect that defining the boundaries of permissible UIP-UDP-trade-offs will be challenging as indicated by literature discussing faithfulness-plausibility trade-offs [11, 14].

Example for property interaction:

"The food was **delicious** and the price somewhat **fair**."

If a user argues for more compactness the explanation could look like this with "fair" not being highlighted:

"The food was **delicious** and the price somewhat fair."

By removing the term "fair" from the feature highlighting the explanation is more compact but less complete.

For works that propose a novel explanation method without a direct use case, it would be helpful to provide information on applicable UIPs and a discussion of how UDPs might impact the performance under which circumstances.

Open Research Directions

The standardization of UIPs and UDPs would enable the development of software-packages that automatically test for applicable properties given an explanation system.

The number of quantitative tests available for UDPs is still low which we attribute to the subjective nature of UDPs. Nevertheless, quantitative tests for UDPs will be helpful and reliable means to better compare methods.

Furthermore, we recommend that newly proposed explanation methods perform ablation studies to simulate possible application scenarios, or at least include discussions to analyze trustworthiness-comprehensibility trade-offs. Likewise, indicating applicable UIPs and reporting on respective statistics should become the standard, similar to the reporting that is customary for predictive performance.

Trust is often stated as goal of explainable ML. Naturally, it would be interesting to investigate how explanation systems developed with the proposed process score with regards to (in-)appropriate trust.

Conclusion

Evaluation of explainable ML is essential to demonstrate trustworthiness and ensure comprehensibility of an explanation system. However, there are no guidelines or standards for evaluation yet. In this work, we argued for quantitative human-grounded evaluation of explainable ML. We presented an evaluation process that includes the user and places importance on quantitative tests. Explanation quality is considered as a construct consisting of several properties. We split the properties into two sets, one to ensure trustworthiness, the other to assess comprehensibility. We highlight that adaptations of the explanation system based on comprehensibility checks can impact properties that quantify trustworthiness. We raise the question whether this trade-off is justifiable under any circumstances and call for a discussion within the community to further develop a position.

Acknowledgements

We thank the reviewers for their constructive feedback and the ML2R Trustworthy ML group and the MLAI Lab for the helpful discussions. This research has been partly funded by the Federal Ministry of Education and Research of Germany as part of the competence center for machine learning ML2R (01|S18038BC). K. Beckh has contributed as part of the SmartHospital.NRW project which is funded by the Ministry for Economic Affairs, Innovation, Digitalization and Energy of the State of North Rhine-Westphalia.

REFERENCES

- [1] Amina Adadi and Mohammed Berrada. 2018. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE access* 6 (2018), 52138–52160.
- [2] David Alvarez-Melis and Tommi S Jaakkola. 2018. On the Robustness of Interpretability Methods. *arXiv*

- preprint arXiv:1806.08049* (2018).
- [3] Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. 2018. Towards better understanding of gradient-based attribution methods for Deep Neural Networks. In *International Conference on Learning Representations*.
- [4] Katharina Beckh, Sebastian Müller, Matthias Jakobs, Vanessa Toborek, Hanxiao Tan, Raphael Fischer, Pascal Welke, Sebastian Houben, and Laura von Rueden. 2021. Explainable Machine Learning with Prior Knowledge: An Overview. *arXiv preprint arXiv:2105.10172* (2021).
- [5] Sebastian Bordt, Michèle Finck, Eric Raidl, and Ulrike von Luxburg. 2022. Post-Hoc Explanations Fail to Achieve their Purpose in Adversarial Contexts. *arXiv preprint arXiv:2201.10295* (2022).
- [6] Nelson Cowan. 2010. The Magical Mystery Four: How Is Working Memory Capacity Limited, and Why? *Current Directions in Psychological Science* 19, 1 (2010), 51–57.
- [7] Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. 2020. A Survey of the State of Explainable AI for Natural Language Processing. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*. ACL, Suzhou, China, 447–459.
- [8] Joseph P. Dietvorst, Berkeley J. Simmons, and Cade Massey. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144, 1 (2015), 114.
- [9] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv preprint arXiv:1702.08608* (2017).
- [10] Upol Ehsan, Q Vera Liao, Michael Muller, Mark O Riedl, and Justin D Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [11] Alon Jacovi and Yoav Goldberg. 2020. Towards Faithfully Interpretable NLP Systems: How Should We Define and Evaluate Faithfulness?. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. ACL, Online, 4198–4205.
- [12] Isaac Lage, Emily Chen, Jeffrey He, Menaka Narayanan, Been Kim, Samuel J Gershman, and Finale Doshi-Velez. 2019. Human Evaluation of Models Built for Interpretability. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 59–67.
- [13] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (April 2020), 1–15.
- [14] Zachary C Lipton. 2018. The Mythos of Model Interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* 16, 3 (2018), 31–57.
- [15] Tim Miller. 2019. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence* 267 (2019), 1–38.

- [16] Sina Mohseni, Jeremy E Block, and Eric Ragan. 2021. Quantitative Evaluation of Machine Learning Explanations: A Human-Grounded Benchmark. In *26th International Conference on Intelligent User Interfaces*. 22–31.
- [17] Christoph Molnar. 2022. *Interpretable machine learning* (second ed.). leanpub.com, Chapter 3.1.
- [18] Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, Jörg Schlötterer, Maurice van Keulen, and Christin Seifert. 2022. From Anecdotal Evidence to Quantitative Evaluation Methods: A Systematic Review on Evaluating Explainable AI. *arXiv preprint arXiv:2201.08164* (2022).
- [19] Andrés Páez. 2019. The Pragmatic Turn in Explainable Artificial Intelligence (XAI). *Minds and Machines* 29, 3 (2019), 441–459.
- [20] Jan L Plass, Roxana Moreno, and Roland Brünken. 2010. Cognitive Load Theory. *Cambridge University Press* (2010).
- [21] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Wortman Vaughan, and Hanna Wallach. 2021. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. ACM.
- [22] Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. ACL*, Online, 4902–4912.
- [23] Mark O Riedl. 2019. Human-Centered Artificial Intelligence and Machine Learning. *Human Behavior and Emerging Technologies* 1, 1 (2019), 33–36.
- [24] Kacper Sokol and Peter Flach. 2020. Explainability Fact Sheets: A Framework for Systematic Assessment of Explainable Approaches. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. ACM, 56–67.
- [25] Giulia Vilone and Luca Longo. 2021. Notions of explainability and evaluation approaches for explainable artificial intelligence. *Information Fusion* 76 (2021), 89–106.
- [26] Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun Suggala, David I Inouye, and Pradeep K Ravikumar. 2019. On the (In)fidelity and Sensitivity of Explanations. In *Advances in Neural Information Processing Systems*, Vol. 32.