

LETTER

Non-algorithms for Explainable Artificial Intelligence[†]

Shane T. Mueller^{*1} | Robert Hoffman² | Gary Klein³ | Tauseef Ibne Mamun¹ | Mohammadreza Jalaieian³

¹Dept. of Cognitive and Learning Sciences,
Michigan Technological University,
Michigan, USA

²MacroCognition, LLC, OH, USA

³Institute for Human and Machine
Cognition, FL, USA

Correspondence

*Corresponding author name,
Corresponding address. Email:

Abstract

The field of Explainable AI (XAI) has focused primarily on algorithms that can help explain decisions and classification and help understand whether a particular action of an AI system is justified. These *XAI algorithms* provide a variety of means for answering a number of questions human users might have about an AI. However, explanation is also supported by *non-algorithms*: methods, tools, interfaces, and evaluations that might help develop or provide explanations for users, either on their own or in company with algorithmic explanations. In this article, we introduce and describe a small number of non-algorithms we have developed. These include several sets of guidelines for methodological guidance about evaluating systems, including both formative and summative evaluation (such as the self-explanation scorecard and stakeholder playbook) and several concepts for generating explanations that can augment or replace algorithmic XAI (such as the Discovery platform, Collaborative XAI, and the Cognitive Tutorial). We will introduce and review several of these example systems, and discuss how they might be useful in developing or improving algorithmic explanations, or even providing complete and useful non-algorithmic explanations of AI and ML systems.

KEYWORDS:

Explainable AI; Human-centered design; Human-AI integration

1 | INTRODUCTION

The typical approach of eXplainable AI (XAI) research starts with two algorithms: one that performs some complex behavior (the AI or ML system) and a second one designed to explain it to users or developers. We refer to this second system as *algorithmic XAI*. Yet for any explanation system, there are likely to be many elements that support the explanation and understanding that are *not* algorithmic¹. Sometimes these are aspects of the interface or architecture²; sometimes these are methods for evaluating the system and determining whether it is effective; sometimes they are tutorials, help documents, and even conversations among

[†]This is an example for title footnote.

⁰**Abbreviations:** XAI, Explainable Artificial Intelligence; ML, Machine Learning; CXAI, Collaborative XAI

developers or users that help them understand the AI. We refer to these, generally, as *non-algorithms*, and the goal of this article is to describe a number of approaches we have developed and implemented that focus on these non-algorithms for XAI. With the extensive work on algorithmic XAI, some discussion of non-algorithmic explanation is useful because (1) substantial aspects of explanation **must** be supported by non-algorithmic approaches; (2) non-algorithms might be valuable on their own for augmenting existing and fielded AI systems without requiring re-engineering with explanation algorithms; and (3) current algorithmic XAI systems may be made more powerful by using non-algorithms.

2 | NON-ALGORITHMS FOR XAI

In Table 1, we identify some example XAI non-algorithms we have developed. One set of methods involves guidance for evaluation, measurement, and validation of XAI systems. Unlike many AI systems which can be evaluated according to performance measures (time, accuracy, efficiency) on an existing data set, true evaluation and validation of explanations need require with human participants and end users.^{3,4,5} Our framework for measuring explanation effectiveness⁶ lays out a suite of measures (including goodness criteria, satisfaction and trust, mental model knowledge, and performance) that can be applied once an AI or XAI system has been developed. However, we have also identified a number of formative evaluation methods that can be applied early during conceptualization and development without requiring *in situ* user evaluation, including measurement according to *explanation goodness*,⁶ *Stakeholder analysis*,⁷ and the *self-explanation scorecard*.² Using these methods, system developers can themselves evaluate their XAI systems.

There are also non-algorithms that can be used to provide explanatory material, adumbrate AI-generated explanations, or support the user's sensemaking or self-explanation effort. These methods include a system we call the *Discovery Platform*, a collaborative approach (*CXAI*) allowing users to explain things and help one another, and a systematic means of developing a *Cognitive Tutorial*, which provides global explanations about a system for novice users.

3 | EXAMPLE SYSTEMS

XAI non-algorithms are motivated by psychological research on self-explanation⁸, sensemaking^{9,10}, naturalistic approaches to explanation¹¹, and other human-centered principles of XAI^{2,7}. One central notion of this approach is that *all explanation involves self-explanation*, in the motivated attempt to develop a satisfactory understanding. This contrasts with what seems to have been the initial premise of the XAI movement that

1. The XAI would generate an explanation,
2. The explanation would be presented to the user,

TABLE 1 Example non-algorithms for supporting explanation and development of XAI

Purpose	Method	Implemented Example
The design of explanations	Mapping explanations to requirements	Stakeholder playbook guidance to tailoring algorithm-generated explanations to needs of different stakeholder roles. Self-explanation scorecard to map XAI-generated explanations on to user's sensemaking requirements
Support for the explaining process	Explanation as exploration	The discovery platform to explore patterns, contrasts, and edge cases
	Explanation as collaboration	CXAI (Collaborative XAI) to allow users to share surprises, provide global explanations, and pose and answer questions
	Global explanation	Cognitive tutorial leverages expert knowledge to provide global explanations and practice exercises
Rigorous experimental evaluation	Methodological guidance	Handbook of experimental design for rigorous assessment of XAI systems with human users
	Data analysis	Methods for determining practical significance of evaluation experiment results

3. The user would understand it, and then

4. Performance, trust, and reliance would improve.

Rather than focusing on a single explanation algorithm, explanatory systems should empower users by giving them information but also supporting interaction that allows them to form and refine explanations they need for their own goals.¹² Next, we will focus on briefly describing the high-level motivations and implementation of several of the novel explanation non-algorithms in Table 1.

3.1 | Stakeholder Playbook

The initial focus of XAI was on explaining AI systems to end-users. The Stakeholder Playbook was created in recognition of the possibility that various “stakeholders” would also need explanations, but also that different stakeholders would need different kinds of explanations, depending on their roles and responsibilities. *The purpose of the Playbook is to enable system developers to appreciate the different ways in which stakeholders need to “look inside” of the AI/XAI system.* For example, some stakeholders, like end-users, might need to understand the boundary conditions of the system (its strengths and limitations). Program managers might need to understand an AI/XAI system in a way that enables them to succinctly explain the system to other people. Leaders of system development teams need to be able to develop appropriate optimism, informed by appropriate skepticism.

While interest in the stakeholder-dependence of explanations has burgeoned in the last few years, there have been only tentative attempts to investigate the matter empirically. We conducted individual cognitive interviews with 18 experienced professionals

concerning their interactions with AI systems. The group including program managers, developers, end-users, legal advocates, and others. Participants were asked just a few questions, including: “What do you feel you need to know about an AI system in order to properly exercise your responsibilities?” and “Can you briefly describe any experiences you have had with AI systems where more knowledge would have helped?”

The interviews resulted in a number of surprises. One of the first surprises we encountered involved the demographics. All the interviewees wore more than one “hat.” A given interviewee might make a comment pertinent from the perspective of the system developer, but then make another comment that pertained to the explanation requirements of an end-user. Thus, it is better to refer to *roles* instead of stakeholder types or groups. That said, we clustered responses according to the following “hats”: jurisprudence specialists, system developers, system development team leaders, procurement or contracting officers, trainers, system evaluators, and policy makers. The answers to the interview questions resulted in a great many discoveries. Here are just three examples:

- Not everyone actually needs or wants an explanation. Only three of the participants spontaneously said that they want explanations of how the AI works. Far more frequent were assertions about the explanation needs of stakeholders’ other than themselves.
- Stakeholders are more likely to need to know about the data than about the AI system that processes the data. Understanding the data the AI system uses would be more helpful than poking under the hood to examine the innards of the system. They wanted to know what data were used to train the AI/ML. They want to know about any system biases.
- Sensemaking by exploration is of greater interest than prepared explanations. A number of participants commented about how they preferred to manipulate (“poke around”) and explore the AI system behavior under different scenarios, to “get a feel for it.” Stakeholders want to be provided with more examples of the AI encountering different situations. End-users said that they would benefit from local explanations that are exploratory rather than discursive: The visualization of trade-offs (e.g., in a scheduling algorithm) would support appropriate reliance and the capacity to anticipate conditions under which anomalous events might occur and the recommendation may be misguided.

For each category we were able to distill explanation requirements. For example, trainers require access to a rich corpus of cases, but especially “edge cases” that allow users to learn. This includes learning how to handle these cases, but also to anticipate when the AI system is entering a brittle zone. For each category we were also able to distill some “cautions.” For example, end-users often require access to the system development team to answer their questions (a requirement) but for end-users, explanation is never a “one-off” — continuing explanation is required as the input data, the work system context, or the operational environment change (a caution). The Stakeholder Playbook itself is a three-page document that can be provided upon

request. The full technical report is also available and provides details of the method and results, including ample quotations from the participants.

3.2 | Self-Explanation Scorecard

Because all users must engage in self-explanation in order to develop an understanding of a system, the Self-Explanation Scorecard² scorecard focuses on aspects of the information people use to explain complex systems to themselves and others. This scorecard is useful for, among other things, early formative evaluation of an explanation concept. An XAI developer can evaluate their envisioned system according to the following criteria:

1. **Features.** Does the XAI highlight importance of features used in a decision?
2. **Successes.** Does the XAI show examples of successful operation?
3. **Mechanisms.** Does the XAI describe mechanisms, rules, or architecture?
4. **AI Reasoning.** Does the XAI provide functional description of algorithms/processes?
5. **Failures.** Does the XAI show examples of failure?
6. **Comparisons.** Does the XAI allow user to compare conditions to draw causal inferences?
7. **Diagnosis of failures.** Does the XAI provide analysis of why failures occur?

The Scorecard helps developers identify the kinds of information that people find useful when generating and refining their own understanding of a complex system². Items earlier in the Scorecard list tend to be simpler, while later items are cognitively more complex and offer deeper insights and require a deeper level of analysis. The intended use of the Scorecard is for developers to self-assess their system early in the design process, and/or to determine whether (sometimes simple) interface or algorithm changes might support higher levels of self-explanation.

The Scorecard can also be used to examine existing systems. To illustrate this, we have coded several published explainable systems according to these criteria, with results that we show in Table 2. Two of the coauthors independently assessed the systems on each dimension of the scorecard. After initial coding, four (out of 42) codes were in disagreement, representing a Cohen's κ of .81. These cases of disagreement (marked with a \pm) were all situations in which both raters thought the explanation was present in the paper, but one rater felt this explanation form was not generated by the XAI system itself, but rather by the authors in support of scientific communication.

This issue was actually more pervasive than just these four cases of disagreement. For many of the explanation types coded as not present in the system, the paper provided information that might be considered supporting the explanation, often by

comparing different versions of models or examining ground truth that might not always be available. However, these were not part of the XAI system itself. *Thus, some algorithms can support numerous “depth of explanation” levels of explanation, but these are often not presented to the system users.* Rather, they are often reserved only the authors’ team and peers. This represented much of the difficulty determining whether a particular level of self-explanation potential characterized some of the systems. Multiple kinds of self-explanation can be supported by an XAI system, but these are not always exposed to the user, and the Scorecard might be useful in identifying alternative ways to present information that might help users.

TABLE 2 Example non-algorithms for supporting explanation and development of XAI

XAI System	1. Features	2. Successes	3. Mechanisms	4. Reasoning	5. Failures	6. Comparisons	7. Diagnoses
Bird classifier XAI ¹³	+	+	-	-	-	±	-
ANN-CBR Twin ¹⁴	-	-	+	±	+	-	-
Partial dependence ¹⁵	+	+	-	-	+	+	-
BaobabView ¹⁶	+	+	+	+	+	-	+
Deconvnet ¹⁷	+	+	+	-	-	-	±
GA ² M ¹⁸	+	+	+	-	-	+	-

Notes: Our understanding is that in subsequent refinements of the model in¹³, additional functions have been explored and implemented as part of the explainable system, including levels 5 and 6.

3.3 | The Discovery platform

This non-Algorithmic Method was developed to support the exploration of the behavior of an AI or XAI system, and thereby satisfy some of the requirements of user self-explanation¹¹. These include:

- *Commonalities and patterns.* Patterns allow user to understand typical cases.
- *Exceptions.* Understanding outliers, anomalies, and exceptions help user isolate and anticipate problematic cases.
- *Failures.* It should be easy to identify errors and mistakes
- *Contrasts.* Contrasting cases allow for easy comparisons, enabling counterfactual and causal reasoning.
- *Confusions.* Identifying high-confusion classes helps anticipate weakness areas of the system.
- *Representations, instances, and examples.* Thumbnails and examples should be visible and browseable.

For an AI system that has a clear training and test corpus—typical for many image classifiers—there will be high-level patterns of performance that cannot be revealed by the usual approach of providing a local justification (e.g., in the form of feature highlighting¹⁹ or a heatmap.²⁰)

The Discovery Platform concept was inspired by conversations we had with XAI system developers who had often browsed hundreds or thousands of image cases of their own data set and discovered systematic problems and strengths with their system, developing special-purpose browsers to help them debug and diagnose their system. Although these tools were developed for internal purposes, they appear to support “explanatory debugging,”^{21,22} for a group of stakeholders, and resemble several existing XAI systems^{16,18} with browseable interfaces that permit exploring data sets and decisions. We considered the functions these support and implemented a general-purpose system that could be used to explore patterns in a variety of image classification data sets.

We implemented a prototype of the Discovery platform using the R web interface *shiny*.¹ The demonstration system uses a simple image support vector machine (SVM) classifier on the MNIST hand-written digit data set²³. We sampled 10,000 test cases to produce a browseable data set on which the system achieved 50% accuracy (the SVM in actuality achieved an 89% accuracy rate). The Discovery Platform provides an interface to select cases based on input class and classifier label, and to sample 5-15 cases based on different simple criteria related to the system’s judged probability across outcome cases. One of the panels in the Discovery Platform is shown in Figure 1.

We have evaluated the discovery platform informally, using three different image classifiers. The system appears to best support elements identifying patterns: common errors; highly confusable images that the system nevertheless identifies accurately, and response biases. Thus, it tends to help generate questions about how the system is working rather than providing justifications. In addition, it has helped us identify and informally test the usefulness of verbal rules that can help a user understand how the system is working (e.g., “when a 4 has a closed top, it is often mistaken for a 9”). It may also be useful coupled with feature-highlighting functions that may provide more causal information about how classifications are made.

3.4 | Example: Collaborative XAI (CXAI)

Social Q&A (SQA) systems²⁴ such as StackExchange have been demonstrated to be a class of non-algorithms that successfully provide explanations for complex software systems. They enable a searchable and browseable database of questions that other users can answer. Those answers can be up-voted to provide social credit to users, and benefit others who have similar questions. We have developed a prototype SQA system focused on XAI which we refer to as *Collaborative XAI* or CXAI, shown in Figure 2), to explore the design and benefits of the explanations that could arise from a social Q&A platform.

¹code available at <https://github.com/stmueller/xai-discovery-platform>

Overview Categories **Contrasts** Cases Data Filtering

Compare two classes and labels

Class1:
1

Label1:
1

Class2:
1

Label2:
7

What to explore?

Random cases
 Worst match to contrast
 Best match to contrast

Contrast Exploration

Probe Case: 1->1 (Random cases)

1->1	1->1	1->1	1->1	1->1
38126 0.96->0.96	55687 0.99->0.99	39404 0.96->0.96	58140 0.98->0.98	57269 0.95->0.95

Probe Case: 1->7 (Random cases)

1->7	1->7	1->7	1->7	1->7
50633 0->0	225 0.01->0.01	50573 0->0	56577 0.27->0.27	34543 0.02->0.02

FIGURE 1 Screenshot of the “contrast explorer” pane of the discovery platform

Although commercial and public SQA sites could be used by XAI developers or users today, they have a number of limitations that may make an XAI-focused SQA platform attractive. The first is that general SQA sites lack of a community-focused mission. Other community-oriented social media platforms (such as discussion forums, discord sites, and subreddits) center on developing a shared community. General SQA sites attempt to expand to a very general audience, and thus lose individual identity and community. Second, general SQA sites are not focused on the known psychological elements of explaining AI. By developing a custom SQA approach, we can incorporate elements that encourage posts based on explanation triggers, and support specific explanatory responses. Finally, general SQA sites are almost necessarily work in isolation. A custom site can be more easily integrated into an AI or XAI interface—with links back to specific cases in a database, and direct links from an explanation interface to existing related posts in the CXAI system, or the ability to easily construct new posts based on cases.

Past uses of SQA suggest that CXAI may be useful in supporting either AI or XAI systems. Research has shown that users of SQA systems are generally satisfied with the information,²⁵ something that is not true for all explanation algorithms. In addition, other features of SQA can help promote better explanations. For example, successful SQA platforms tend to develop an engaged community and encourage participation via up-votes and bounties,²⁶ and the best answers in a SQA platform have been shown to correlate with the consistent participation of users.²⁷ Next, one may worry that a collaborative explanation system will suffer from having no authoritative answers. Yet successful SQA sites do not generally enlist professional or verified expert providing

CrowdCollaboration

Search....

If you use a search suggestion for search, that search will find entries whose tag/title/topic/author name will match your search term exactly. You can always disregard the suggestions that will be shown to you and search entries using a word/string of words, this type of search will bring back search results that have your word/string of words in the results

Use definite keyword/keywords for better search result

Artificial Intelligent systems are not perfect. They do not really think like people. They may do things that surprise you. Using CrowdCollaboration you can leave behind bread crumbs or signposts. Your postings will help other people understand and work with the Image Classification System.

This collaborative effort is for a/an

Image Classification System
AI Web

Possible Topics

HOW IT WORKS

1. What does it achieve?
2. What can't it do?

SURPRISES and MYSTERIES

1. Why did it do that?
2. Why didn't it do x?

TRICKS & DISCOVERIES

1. Here's something that surprised me.
2. Here's a trick I discovered.

TRAPS

1. What do I have to look out for?
2. What do I do if it gets something wrong?
3. How can it fool me?
4. What do I do if I do not trust what it did?

Post(s)

Plier as nipper
lubuntu

A common label, 'nipper' is used for pliers in the original transformation if the plier is not in its package. I think the AI gets confused with the packaging. But it gives the label 'pliers' for all the pliers in the original transformation.

11/20/2020 13:41:34 0 Answer(s)

Rotate ransform only causes problems when there is a square background
shane

When images have a white background, the rotate transform does not usually change the accuracy, because the tools are usually at an arbitrary angle anyway. When there is background so the rotated image produces a diamond-shape, the AI is often confused and says 'angle'

11/13/2020 12:49:35 0 Answer(s)

Stats

Posts (58)

Answers (24)

Recent Tags

flashlight non-target object

tape measure

measuring instrument

color frame

FIGURE 2 Screenshot of the CXAI system

good answers, and allow users to build a reputation within a particular question category and become known as an expert on the site.²⁴ They also allow users to vote on answers to collaboratively verify they are accurate. Thus, good explanations may surface without the help of an authoritative expert. Another promising aspect of SQA is that, with a proper incentives and an engaged user base, answers emerge quickly. One study²⁸ showed that in an aggregate sample of over 3 million questions and 16 million answers in Yahoo! Answers, 30% of questions received an answer within 5 minutes of submission and 92% received an answer within an hour.

3.5 | Example: Cognitive Tutorials for AI and XAI

Researchers in the field of XAI have been discussing the distinction between global and local explanation from the outset (although the distinction can be traced to earlier work on causal reasoning). This distinction covers the *scope* or *focus* of the

explanation—with “global” describing how the system architecture works in general (perhaps covering algorithms, training materials, specific sets of rules, and patterns of behavior), and “local” describing how a particular case was handled. Modern XAI algorithms primarily deal with local explanations (or *justifications*). These make sense to computer scientists, and explain to them why the system developers “did it that way.” They are also useful for trying to understand issues of fairness, justice, and to identify remedies: a local explanation of why a loan was denied will hopefully let a developer understand if the denial stemmed from something improper, and will help an applicant determine what they can do to improve their chances of approval. However, sole reliance on local explanations is perhaps at odds with developing long-term general trust and understanding in the system, because local explanations are myopic and analytical.

One method we have explored for producing global explanations is what we call a Cognitive Tutorial, which uses experiential training in the form of a user guide about the cognitive operations of the AI^{29,30}. The Cognitive Tutorial recognizes that users will come to the AI with misconceptions about how it works—often assuming it works in the same way a human would³¹. However, AI often succeeds and fails in unexpected ways. The goal of the cognitive tutorial is to use experiential training to help the user understand the competence boundaries of the system—along dimensions that include modeling/representation, algorithms, data, and output/visualization.

We have developed a Cognitive Tutorial Authoring Guide, available from the authors on request. This guide details the specific steps and procedures for identifying learning objectives and implementing several kinds of learning modules, including: How to Use It, How to Use It Improperly, Common Misconceptions, Novel Problems, Forced-choice decisions, and Decision rule learning. Critical pedagogical methods we explore include showing the user both success and failure modes, creating a “counterfactual sequence” that illustrates how a increasing small change to a larger change in input features will change the outcome, and the use of examples to reinforce and learn rules of thumb about AI system performance.

We believe that a tutorial can provide an important initial exposure to an AI system. Although it involves extensive human involvement, it may ultimately provide a better or faster learning path than typical algorithmic approaches that do not attempt to teach the user general patterns, and instead focus on justifying specific decisions. We have found that many approaches explored in algorithmic XAI can be useful in helping design and implement tutorial-based training.

4 | CONCLUSION

The combination of Algorithmic and Non-algorithmic approaches to XAI is likely to be far more successful than any purely Algorithmic approach. In general, our non-algorithms do not replace algorithms, but complement them and help provide missing information, help suggest gaps and missing information, and help communicate high-level actionable information to users.

This derives from the psychology of explanation: explanation as an exploratory and collaborative process for ensuring that AI technology is part of an understandable, learnable, usable, and useful human-machine work system.

Acknowledgements

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). Acknowledgement and Disclaimer: This material is approved for public release. Distribution is unlimited. This material is based on research sponsored by the Air Force Research Lab (AFRL) under agreement number FA8650-17-2-7711. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL or the U.S. Government. The authors have no conflicts of interest to report with respect to the work described here.

References

1. VanLehn K, Ball W, Kowalski B. Explanation-based learning of correctness: Towards a model of the self-explanation effect. tech. rep., CARNEGIE-MELLON UNIV PITTSBURGH PA DEPT OF PSYCHOLOGY; : 1990.
2. Mueller ST, Veinott ES, Hoffman RR, et al. Principles of Explanation in Human-AI Systems. *arXiv preprint arXiv:2102.04972* 2021.
3. Miller T, Howe P, Sonenberg L. Explainable AI: Beware of Inmates Running the Asylum. *IJCAI-17 Workshop on Explainable Artificial Intelligence (XAI)* 2017.
4. Doshi-Velez F, Kim B. A Roadmap for a Rigorous Science of Interpretability. *arXiv preprint arXiv:1702.08608* 2017.
5. Doshi-Velez F, Kim B. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* 2017.
6. Hoffman RR, Mueller ST, Klein G, Litman J. Metrics for explainable AI: Challenges and prospects. *arXiv preprint arXiv:1812.04608* 2018.
7. Clancey WJ, Hoffman RR. Methods and standards for research on explainable artificial intelligence: Lessons from intelligent tutoring systems research. tech. rep., DARPA; Alexandria, VA: 2021.
8. Chi MT, Bassok M, Lewis MW, Reimann P, Glaser R. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science* 1989; 13(2): 145–182.

9. Klein G, Moon B, Hoffman RR. Making sense of sensemaking 1: Alternative perspectives. *IEEE intelligent systems* 2006; 21(4): 70–73.
10. Klein GA, Moon B, Hoffman RR. Making sense of sensemaking 2: A macrocognitive model. *IEEE intelligent systems* 2006; 21(5): 88–92. 00404.
11. Klein G, Hoffman R, Mueller S. Naturalistic psychological model of explanatory reasoning: How people explain things to others and to themselves. In: *International Conference on Naturalistic Decision Making*; 2019.
12. Brézillon P. Context needs in cooperative building of explanations. In: *First European Conference on Cognitive Science in Industry*; 1994: 443–450. 00024.
13. Hendricks LA, Akata Z, Rohrbach M, Donahue J, Schiele B, Darrell T. Generating visual explanations. In: *European Conference on Computer Vision*Springer; 2016: 3–19. 00016.
14. Ford C, Kenny EM, Keane MT. Play MNIST For Me! User Studies on the Effects of Post-Hoc, Example-Based Explanations & Error Rates on Debugging a Deep Learning, Black-Box Classifier. *arXiv preprint arXiv:2009.06349* 2020.
15. Krause J, Perer A, Bertini E. Using visual analytics to interpret predictive machine learning models. *arXiv preprint arXiv:1606.05685* 2016.
16. Van Den Elzen S, Van Wijk JJ. Baobabview: Interactive construction and analysis of decision trees. In: *2011 IEEE conference on visual analytics science and technology (VAST)*IEEE; 2011: 151–160.
17. Zeiler MD, Fergus R. Visualizing and understanding convolutional networks. In: *European conference on computer vision*Springer; 2014: 818–833. 01903.
18. Caruana R, Lou Y, Gehrke J, Koch P, Sturm M, Elhadad N. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In: *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*; 2015: 1721–1730.
19. Ribeiro M, Singh S, Guestrin C. Model-Agnostic Interpretability of Machine Learning. In: *Proceedings of the 2016 ICML Workshop on Interpretability in Machine Learning (WHI 2016)*; 2016; Ney York, NY.
20. Selvaraju RR, Das A, Vedantam R, Cogswell M, Parikh D, Batra D. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *arXiv preprint arXiv:1610.02391* 2016.
21. Kulesza T, Stumpf S, Burnett M, et al. Explanatory debugging: Supporting end-user debugging of machine-learned programs. In: *Visual Languages and Human-Centric Computing (VL/HCC), 2010 IEEE Symposium on*IEEE; 2010: 41–48.

22. Kulesza T, Burnett M, Wong WK, Stumpf S. Principles of explanatory debugging to personalize interactive machine learning. In: *Proceedings of the 20th International Conference on Intelligent User Interfaces*; 2015; Atlanta, Georgia: 126–137
23. Deng L. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 2012; 29(6): 141–142. ISBN: 1053-5888 Publisher: IEEE.
24. Shah C, Oh S, Oh JS. Research agenda for social Q&A. *Library & Information Science Research* 2009; 31(4): 205–209. ISBN: 0740-8188 Publisher: Elsevier.
25. Gazan R. Social Q&A. *Journal of the American Society for Information Science and Technology* 2011; 62(12): 2301–2312. ISBN: 1532-2882 Publisher: Wiley Online Librarydoi: <https://doi.org/10.1002/asi.21562>
26. Zhou J, Wang S, Bezemer CP, Hassan AE. Bounties on technical Q&A sites: a case study of Stack Overflow bounties. *Empirical Software Engineering* 2020; 25(1): 139–177. ISBN: 1573-7616 Publisher: Springer.
27. Nam KK, Ackerman MS, Adamic LA. Questions in, knowledge in? A study of Naver’s question answering community. *Proceedings of the SIGCHI conference on human factors in computing systems* 2009: 779–788.
28. Shah C. Measuring effectiveness and user satisfaction in Yahoo! Answers. *First Monday* 2011; 16(2). doi: <https://doi.org/10.5210/fm.v16i2.3092>
29. Mueller ST, Klein G, Burns C. Experiencing the tool without experiencing the pain: concepts for an experiential user guide. In: *Proc. 9th Int’l Conf. Naturalistic Decision Making*; 2009: 105–112.
30. Mueller S, Tan S, Klein G. How many personality configurations exist?. Psych Today Blog: Seeing what others don’t; 2019.
31. Mueller ST. Cognitive anthropomorphism of AI: How Humans and Computers Classify images. *Ergonomics in Design* 2020; 28(3): 12–19. Publisher: SAGE Publications.

How to cite this article: Mueller, S. T., Klein, G., Hoffman, R. R., and T. Mamun (2021), Non-algorithms for explanation in AI Systems, *Applied AI Letters*, 2021.