

Explanation-Based Human Debugging of NLP Models: A Survey

Piyawat Lertvittayakumjorn and Francesca Toni

Department of Computing
Imperial College London, UK
{p11515, ft}@imperial.ac.uk

Abstract

Debugging a machine learning model is hard since the bug usually involves the training data and the learning process. This becomes even harder for an opaque deep learning model if we have no clue about how the model actually works. In this survey, we review papers that exploit explanations to enable humans to give feedback and debug NLP models. We call this problem *explanation-based human debugging (EBHD)*. In particular, we categorize and discuss existing work along three dimensions of EBHD (the bug context, the workflow, and the experimental setting), compile findings on how EBHD components affect the feedback providers, and highlight open problems that could be future research directions.

1 Introduction

Explainable AI focuses on generating explanations for AI models as well as for their predictions. It is gaining more and more attention these days since explanations are necessary in many applications, especially in high-stake domains such as healthcare, law, transportation, and finance (Adadi and Berrada, 2018). Some researchers have explored various merits of explanations to humans, such as supporting human decision making (Lai and Tan, 2019; Lertvittayakumjorn et al., 2021), increasing human trust in AI (Jacovi et al., 2020), and even teaching humans to perform challenging tasks (Lai et al., 2020). On the other hand, explanations can benefit the AI systems as well, for example, when explanations are used to promote system acceptance (Cramer et al., 2008), to verify the model reasoning (Caruana et al., 2015), and to find potential causes of errors (Han et al., 2020).

In this paper, we review progress to date specifically on how explanations have been used in the literature to enable humans to fix bugs in NLP models. We refer to this research area as *explanation-based human debugging (EBHD)*, as

a general umbrella term encompassing *explanatory debugging* (Kulesza et al., 2010) and *human-in-the-loop debugging* (Lertvittayakumjorn et al., 2020). We define EBHD as the process of fixing or mitigating bugs in a trained model using human feedback given in response to explanations for the model. EBHD is helpful when the training data at hand leads to suboptimal models (due, for instance, to biases or artifacts in the data), and hence human knowledge is needed to verify and improve the trained models. In fact, EBHD is related to three challenging and intertwined issues in NLP: explainability (Danilevsky et al., 2020), interactive and human-in-the-loop learning (Amershi et al., 2014; Wang et al., 2021), and knowledge integration (von Rueden et al., 2021; Kim et al., 2021). Although there are overviews for each of these topics (as cited above), our paper is the first to draw connections among the three towards the final application of model debugging in NLP.

Whereas most people agree on the meaning of the term *bug* in software engineering, various meanings have been ascribed to this term in machine learning (ML) research. For example, Selsam et al. (2017) considered bugs as implementation errors, similar to software bugs, while Cadamuro et al. (2016) defined a bug as a particularly damaging or inexplicable test error. In this paper, we follow the definition of (model) bugs from Adebayo et al. (2020) as contamination in the learning and/or prediction pipeline that makes the model produce incorrect predictions or learn error-causing associations. Examples of bugs include spurious correlation, labeling errors, and undesirable behavior in out-of-distribution testing.

The term *debugging* is also interpreted differently by different researchers. Some consider debugging as a process of identifying or uncovering causes of model errors (Parikh and Zitnick, 2011; Graliński et al., 2019), while others stress that debugging must not only reveal the causes of problems but also fix or mitigate them (Kulesza et al.,

Paper	Context			Workflow				Setting
	Task	Model	Bug sources	Exp. scope	Exp. method	Feedback	Update	
Kulesza et al. (2009)	TC	NB	AR	G,L	SE	LB,WS	M,D	SP
Stumpf et al. (2009)	TC	NB	SS	L	SE	WO	T	SP
Kulesza et al. (2010)	TC	NB	SS	G,L	SE	WO,LB	M,D	SP
Kulesza et al. (2015)	TC	NB	AR	G,L	SE	WO,WS	M	SP
Ribeiro et al. (2016)	TC	SVM	AR	L	PH	WO	D	CS
Koh and Liang (2017)	TC	LR	WL	L	PH	LB	D	SM
Ribeiro et al. (2018b)	VQA	TellQA	AR	G	PH	RU	D	SP
	TC	fastText	AR,OD					
Teso and Kersting (2019)	TC	LR	AR	L	PH	WO	D	SM
Cho et al. (2019)	TQA	NeOp	AR	L	SE	AT	T	NR
Khanna et al. (2019)	TC	LR	WL	L	PH	LB	D	SM
Lertvittayakumjorn et al. (2020)	TC	CNN	AR,SS,OD	G	PH	FE	T	CS
Smith-Renner et al. (2020)	TC	NB	AR,SS	L	SE	LB,WO	M,D	CS
Han and Ghosh (2020)	TC	LR	WL	L	PH	LB	D	SM
Yao et al. (2021)	TC	BERT*	AR,OD	L	PH	RE	D,T	SP
Zylberajch et al. (2021)	NLI	BERT	AR	L	PH	ES	D	SP

Table 1: Overview of existing work on EBHD of NLP models. We use abbreviations as follows: **Task**: TC = Text Classification (single input), VQA = Visual Question Answering, TQA = Table Question Answering, NLI = Natural Language Inference / **Model**: NB = Naive Bayes, SVM = Support Vector Machines, LR = Logistic Regression, TellQA = Telling QA, NeOp = Neural Operator, CNN = Convolutional Neural Networks, BERT* = BERT and RoBERTa / **Bug sources**: AR = Natural artifacts, SS = Small training subset, WL = Wrong label injection, OD = Out-of-distribution tests / **Exp. scope**: G = Global explanations, L = Local explanations / **Exp. method**: SE = Self-explaining, PH = Post-hoc method / **Feedback** (form): LB = Label, WO = Word(s), WS = Word(s) Score, ES = Example Score, FE = Feature, RU = Rule, AT = Attention, RE = Reasoning / **Update**: M = Adjust the model parameters, D = Improve the training data, T = Influence the training process / **Setting**: SP = Selected participants, CS = Crowdsourced participants, SM = Simulation, NR = Not reported.

2015; Yousefzadeh and O’Leary, 2019). In this paper, we adopt the latter interpretation.

Scope of the Survey. We focus on work using explanations of NLP models to expose whether there are bugs and exploit human feedback to fix the bugs (if any). To collect relevant papers, we started from some pivotal EBHD work (e.g., Kulesza et al., 2015; Ribeiro et al., 2016; Teso and Kersting, 2019), and added EBHD papers citing or being cited by the pivotal work (e.g., Stumpf et al., 2009; Kulesza et al., 2010; Lertvittayakumjorn et al., 2020; Yao et al., 2021). Next, to ensure that we did not miss any important work, we searched for papers on Semantic Scholar¹ using the Cartesian product of five keyword sets: {debugging}, {text, NLP}, {human, user, interactive, feedback}, {explanation, explanatory}, and {learning}. With 16 queries in total, we collected the top 100 pa-

pers (ranked by relevancy) for each query and kept only the ones appearing in at least 2 out of the 16 query results. This resulted in 234 papers that we then manually checked, leading to selecting a few additional papers, including Han and Ghosh (2020) and Zylberajch et al. (2021). The overall process resulted in 15 papers listed in Table 1 as the *selected studies* primarily discussed in this survey. In contrast, some papers from the following categories appeared in the search results, but were not selected because, strictly speaking, they are not in the main scope of this survey: debugging without explanations (Kang et al., 2018), debugging outside the NLP domain (Ghai et al., 2021; Popordanoska et al., 2020; Bekkemoen and Langseth, 2021), refining the ML pipeline instead of the model (Lourenço et al., 2020; Schoop et al., 2020), improving the explanations instead of the model (Ming et al., 2019), and work centered on revealing but not fixing bugs (Ribeiro et al., 2020; Krause et al., 2016; Krishnan and Wu, 2017).

¹<https://www.semanticscholar.org/>.

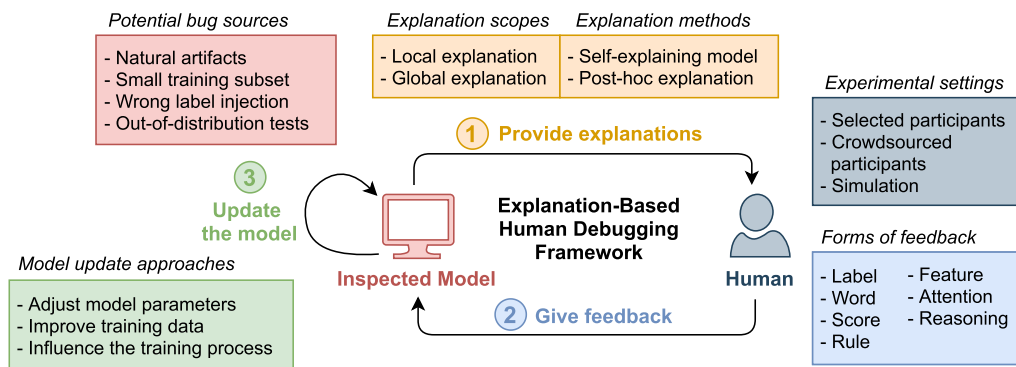


Figure 1: A general framework for explanation-based human debugging (EBHD) of NLP models, consisting of the inspected (potentially buggy) model, the humans providing feedback, and a three-step workflow. Boxes list examples of the options (considered in the selected studies) for the components or steps in the general framework.

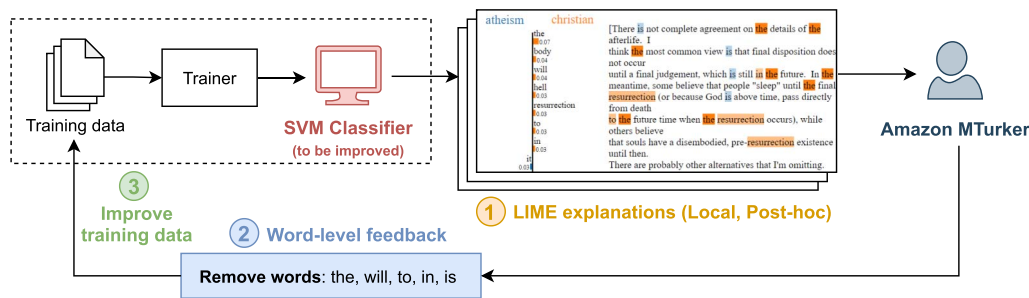


Figure 2: The proposal by Ribeiro et al. (2016) as an instance of the general EBHD framework.

General Framework. EBHD consists of three main steps as shown in Figure 1. First, the explanations, which provide interpretable insights into the inspected model and possibly reveal bugs, are given to humans. Then, the humans inspect the explanations and give feedback in response. Finally, the feedback is used to update and improve the model. These steps can be carried out once, as a one-off improvement, or iteratively, depending on how the debugging framework is designed.

As a concrete example, Figure 2 illustrates how Ribeiro et al. (2016) improved an SVM text classifier trained on the 20Newsgroups dataset (Lang, 1995). This dataset has many artifacts that could make the model rely on wrong words or tokens when making predictions, reducing its generalizability.² To perform EBHD, Ribeiro et al. (2016) recruited humans from a crowdsourcing platform (i.e., Amazon Mechanical Turk) and asked them to inspect LIME explanations³ (i.e., word relevance

scores) for model predictions of ten examples. Then, the humans gave feedback by identifying words in the explanations that should not have received high relevance scores (i.e., supposed to be the artifacts). These words were then removed from the training data, and the model was re-trained. The process was repeated for three rounds, and the results show that the model generalized better after every round. Using the general framework in Figure 1, we can break the framework of Ribeiro et al. (2016) into components as depicted in Figure 2. Throughout the paper, when reviewing the selected studies, we will use the general framework in Figure 1 for analysis, comparison, and discussion.

Human Roles. To avoid confusion, it is worth noting that there are actually two human roles in the EBHD process. One, of course, is that of *feedback provider(s)*, looking at the explanations and providing feedback (noted as ‘Human’ in Figure 1). The other is that of *model developer(s)*, training the model and organizing the EBHD process (not shown in Figure 1). In practice, a person could be both model developer

²For more details, please see Section 2.1.3.

³LIME stands for Local Interpretable Model agnostic Explanations (Ribeiro et al., 2016). For each model prediction, it returns relevance scores for words in the input text to show how important each word is for the prediction.

and feedback provider. This usually happens during the model validation and improvement phase, where the developers try to fix the bugs themselves. Sometimes, however, other stakeholders could also take the feedback provider role. For instance, if the model is trained to classify electronic medical records, the developers (who are mostly ML experts) hardly have the medical knowledge to provide feedback. So, they may ask doctors acting as consultants to the development team to be the feedback providers during the model improvement phase. Further, EBHD can be carried out after deployment, with end users as the feedback providers. For example, a model auto-suggesting the categories of new emails to end users can provide explanations supporting the suggestions as part of its normal operation. Also, it can allow the users to provide feedback to both the suggestions and the explanations. Then, a routine written by the developers will be triggered to process the feedback and update the model automatically to complete the EBHD workflow. In this case, we need to care about the trust, frustration, and expectation of the end users while and after they give feedback. In conclusion, EBHD can take place practically both before and after the model is deployed, and many stakeholders can act as the feedback providers, including, but not limited to, the model developers, the domain experts, and the end users.

Paper Organization. Section 2 explains the choices made by existing work to achieve EBHD of NLP models. This illustrates the current state of the field with the strengths and limitations of existing work. Naturally, though, a successful EBHD framework cannot neglect the “imperfect” nature of feedback providers, who may not be an ideal oracle. Hence, Section 3 compiles relevant human factors that could affect the effectiveness of the debugging process as well as the satisfaction of the feedback providers. After that, we identify open challenges of EBHD for NLP in Section 4 before concluding the paper in Section 5.

2 Categorization of Existing Work

Table 1 summarizes the selected studies along three dimensions, amounting to the debugging context (i.e., tasks, models, and bug sources), the workflow (i.e., the three steps in our general framework), and the experimental setting (i.e., the mode

of human engagement). We will discuss these dimensions with respect to the broader knowledge of explainable NLP and human-in-the-loop learning, to shed light on the current state of EBHD of NLP models.

2.1 Context

To demonstrate the debugging process, existing work needs to set up the bug situation they aim to fix, including the target NLP task, the inspected ML model, and the source of the bug to be addressed.

2.1.1 Tasks

Most papers in Table 1 focus on text classification with single input (TC) for a variety of specific problems such as email categorization (Stumpf et al., 2009), topic classification (Kulesza et al., 2015; Teso and Kersting, 2019), spam classification (Koh and Liang, 2017), sentiment analysis (Ribeiro et al., 2018b), and auto-coding of transcripts (Kulesza et al., 2010). By contrast, Zylberajch et al. (2021) targeted natural language inference (NLI) which is a type of text-pair classification, predicting whether a given premise entails a given hypothesis. Finally, two papers involve question answering (QA), i.e., Ribeiro et al. (2018b) (focusing on visual question answering [VQA]) and Cho et al. (2019) (focusing on table question answering [TQA]).

Ghai et al. (2021) suggested that most researchers work on TC because, for this task, it is much easier for lay participants to understand explanations and give feedback (e.g., which keywords should be added or removed from the list of top features).⁴ Meanwhile, some other NLP tasks require the feedback providers to have linguistic knowledge such as part-of-speech tagging, parsing, and machine translation. The need for linguists or experts renders experiments for these tasks more difficult and costly. However, we suggest that there are several tasks where the trained models are prone to be buggy but the tasks are underexplored in the EBHD setting, though they are not too difficult to experiment on with lay people. *NLI*, the focus of Zylberajch et al. (2021), is one of them. Indeed, McCoy et al. (2019) and Gururangan et al. (2018) showed that NLI models

⁴Nevertheless, some specific TC tasks, such as authorship attribution (Juola, 2007) and deceptive review detection (Lai et al., 2020), are exceptions because lay people are generally not good at these tasks. Thus, they are not suitable for EBHD.

can exploit annotation artifacts and fallible syntactic heuristics to make predictions rather than learning the logic of the actual task. Other tasks and their bugs include: *QA*, where Ribeiro et al. (2019) found that the answers from models are sometimes inconsistent (i.e., contradicting previous answers); and *reading comprehension*, where Jia and Liang (2017) showed that models, which answer a question by reading a given paragraph, can be fooled by an irrelevant sentence being appended to the paragraph. These non-TC NLP tasks would be worth exploring further in the EBHD setting.

2.1.2 Models

Early work used Naive Bayes models with bag-of-words (NB) as text classifiers (Kulesza et al., 2009, 2010; Stumpf et al., 2009), which are relatively easy to generate explanations for and to incorporate human feedback into (discussed in Section 2.2). Other traditional models used include logistic regression (LR) (Teso and Kersting, 2019; Han and Ghosh, 2020) and support vector machines (SVM) (Ribeiro et al., 2016), both with bag-of-words features. The next generation of tested models involves word embeddings. For text classification, Lertvittayakumjorn et al. (2020) focused on convolutional neural networks (CNN) (Kim, 2014) and touched upon bidirectional LSTM networks (Hochreiter and Schmidhuber, 1997), while Ribeiro et al. (2018b) used fastText, relying also on n -gram features (Joulin et al., 2017). For VQA and TQA, the inspected models used attention mechanisms for attending to relevant parts of the input image or table. These models are Telling QA (Zhu et al., 2016) and Neural Operator (NeOp) (Cho et al., 2018), used by Ribeiro et al. (2018b) and Cho et al. (2019), respectively. While the NLP community nowadays is mainly driven by pre-trained language models (Qiu et al., 2020) with many papers studying their behaviors (Rogers et al., 2021; Hoover et al., 2020), only Zylberajch et al. (2021) and Yao et al. (2021) have used pre-trained language models, including BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), as test beds for EBHD.

2.1.3 Bug Sources

Most of the papers in Table 1 experimented on training datasets with natural artifacts (AR), which cause spurious correlation bugs (i.e., the input texts having signals that are correlated to but not

the reasons for specific outputs) and undermine models' generalizability. Out of the 15 papers we surveyed, 5 used the 20Newsgroups dataset (Lang, 1995) as a case study, because it has many natural artifacts. For example, some punctuation marks appear more often in one class due to the writing styles of the authors contributing to the class, so the model uses these punctuation marks as clues to make predictions. However, because 20Newsgroups is a topic classification dataset, a better model should focus more on the topic of the content since the punctuation marks can also appear in other classes, especially when we apply the model to texts in the wild. Apart from classification performance drops, natural artifacts can also cause model biases, as shown in De-Arteaga et al. (2019) and Park et al. (2018) and debugged in Lertvittayakumjorn et al. (2020) and Yao et al. (2021).

In the absence of strong natural artifacts, bugs can still be simulated using several techniques. First, using only a small subset of labeled data (SS) for training could cause the model to exploit spurious correlation leading to poor performance (Kulesza et al., 2010). Second, injecting wrong labels (WL) into the training data can obviously blunt the model quality (Koh and Liang, 2017). Third, using out-of-distribution tests (OD) can reveal that the model does not work effectively in the domains that it has not been trained on (Lertvittayakumjorn et al., 2020; Yao et al., 2021). All of these techniques give rise to undesirable model behaviors, requiring debugging. Another technique, not found in Table 1 but suggested in related work (Idahl et al., 2021), is contaminating input texts in the training data with decoys (i.e., injected artifacts) which could deceive the model into predicting for the wrong reasons. This has been experimented with in the computer vision domain (Rieger et al., 2020), and its use in the EBHD setting in NLP could be an interesting direction to explore.

2.2 Workflow

This Section describes existing work around the three steps of the EBHD workflow in Figure 1, namely, how to generate and present the explanations, how to collect human feedback, and how to update the model using the feedback. Researchers need to make decisions on these key points harmoniously to create an effective debugging workflow.

2.2.1 Providing Explanations

The main role of explanations here is to provide interpretable insights into the model and uncover its potential misbehavior or irrationality, which sometimes cannot be noticed by looking at the model outputs or the evaluation metrics.

Explanation Scopes. Basically, there are two main types of explanations that could be provided to feedback providers. Local explanations (L) explain the predictions by the model for individual inputs. In contrast, global explanations (G) explain the model overall, independently of any specific inputs. It can be seen from Table 1 that most existing work use local explanations. One reason for this may be that, for complex models, global explanations can hardly reveal details of the models' inner workings in a comprehensible way to users. So, some bugs are imperceptible in such high-level global explanations and then not corrected by the users. For example, the debugging framework *FIND*, proposed by Lertvittayakumjorn et al. (2020), uses only global explanations, and it was shown to work more effectively on significant bugs (such as gender bias in abusive language detection) than on less-obvious bugs (such as dataset shift between product types of sentiment analysis on product reviews). Otherwise, Ribeiro et al. (2018b) presented adversarial replacement rules as global explanations to reveal the model weaknesses only, without explaining how the whole model worked.

On the other hand, using local explanations has limitations in that it demands a large amount of effort from feedback providers to inspect the explanation of every single example in the training/validation set. With limited human resources, efficient ways to rank or select examples to explain would be required (Idahl et al., 2021). For instance, Khanna et al. (2019) and Han and Ghosh (2020) targeted explanations of incorrect predictions in the validation set. Ribeiro et al. (2016) picked sets of non-redundant local explanations to illustrate the global picture of the model. Instead, Teso and Kersting (2019) leveraged heuristics from active learning to choose unlabeled examples that maximize some informativeness criteria.

Recently, some work in explainable AI considers generating explanations for a group of predictions (Johnson et al., 2020; Chan et al., 2020) (e.g., for all the false positives of a certain class),

thus staying in the middle of the two extreme explanation types (i.e., local and global). This kind of explanation is not too fine-grained, yet it can capture some suspicious model behaviors if we target the right group of examples. So, it would be worth studying in the context of EBHD (to the best of our knowledge, no existing study experiments with it).

Generating Explanations. To generate explanations in general, there are two important questions we need to answer. First, which format should the explanations have? Second, how do we generate the explanations?

For the first question, we see many possible answers in the literature of explainable NLP (e.g., see the survey by Danilevsky et al., 2020). For instance, *input-based explanations* (so-called feature importance explanations) identify parts of the input that are important for the prediction. The explanation could be a list of importance scores of words in the input, so-called *attribution scores* or *relevance scores* (Lundberg and Lee, 2017; Arras et al., 2016). *Example-based explanations* select influential, important, or similar examples from the training set to explain why the model makes a specific prediction (Han et al., 2020; Guo et al., 2020). *Rule-based explanations* provide interpretable decision rules that approximate the prediction process (Ribeiro et al., 2018a). *Adversarial-based explanations* return the smallest changes in the inputs that could change the predictions, revealing the model misbehavior (Zhang et al., 2020a). In most NLP tasks, input-based explanations are the most popular approach for explaining predictions (Bhatt et al., 2020). This is also the case for EBHD, as most selected studies use input-based explanations (Kulesza et al., 2009, 2010; Teso and Kersting, 2019; Cho et al., 2019) followed by example-based explanations (Koh and Liang, 2017; Khanna et al., 2019; Han and Ghosh, 2020). Meanwhile, only Ribeiro et al. (2018b) use adversarial-based explanations, whereas Stumpf et al. (2009) experiment with input-based, rule-based, and example-based explanations.

For the second question, there are two ways to generate the explanations: self-explaining methods and post-hoc explanation methods. Some models (e.g., Naive Bayes, logistic regression, and decision trees) are *self-explaining* (SE) (Danilevsky et al., 2020), also referred to as transparent

(Adadi and Berrada, 2018) or inherently interpretable (Rudin, 2019). Local explanations of self-explaining models can be obtained at the same time as predictions, usually from the process of making those predictions, while the models themselves can often serve directly as global explanations. For example, feature importance explanations for a Naive Bayes model can be directly derived from the likelihood terms in the Naive Bayes equation, as done by several papers in Table 1 (Kulesza et al., 2009; Smith-Renner et al., 2020). Also, using attention scores on input as explanations, as done in Cho et al. (2019), is a self-explaining method because the scores were obtained during the prediction process.

In contrast, *post-hoc explanation methods* (PH) perform additional steps to extract explanations after the model is trained (for a global explanation) or after the prediction is made (for a local explanation). If the method is allowed to access model parameters, it may calculate word relevance scores by propagating the output scores back to the input words (Arras et al., 2016) or analyzing the derivative of the output with respect to the input words (Smilkov et al., 2017; Sundararajan et al., 2017). If the method cannot access the model parameters, it may perturb the input and see how the output changes to estimate the importance of the altered parts of the input (Ribeiro et al., 2016; Jin et al., 2020). The important words and/or the relevance scores can be presented to the feedback providers in the EBHD workflow in many forms such as a list of words and their scores (Teso and Kersting, 2019; Ribeiro et al., 2016), word clouds (Lertvittayakumjorn et al., 2020), and a parse tree (Yao et al., 2021). Meanwhile, the influence functions method, used in Koh and Liang (2017) and Zylberajch et al. (2021), identifies training examples which influence the prediction by analyzing how the prediction would change if we did not have each training point. This is another post-hoc explanation method as it takes place after prediction. It is similar to the other two example-based explanation methods used in (Khanna et al., 2019; Han and Ghosh, 2020).

Presenting Explanations. It is important to carefully design the presentation of explanations, taking into consideration the background knowledge, desires, and limits of the feedback providers. In the debugging application by Kulesza et al. (2009), lay users were asked to provide feedback

to email categorizations predicted by the system. The users were allowed to ask several Why questions (inspired by Myers et al., 2006) through either the menu bar, or by right-clicking on the object of interest (such as a particular word). Examples include “Why will this message be filed to folder A?”, “Why does word x matter to folder B?”. The system then responded by textual explanations (generated using templates), together with visual explanations such as bar plots for some types of questions. All of these made the interface become more user-friendly. In 2015, Kulesza et al. proposed, as desirable principles, that the presented explanations should be sound (i.e., truthful in describing the underlying model), complete (i.e., not omitting important information about the model), but not overwhelming (i.e., remaining comprehensible). However, these principles are challenging especially when working on non-interpretable complex models.

2.2.2 Collecting Feedback

After seeing explanations, humans generally desire to improve the model by giving feedback (Smith-Renner et al., 2020). Some existing work asked humans to confirm or correct machine-computed explanations. Hence, the form of feedback fairly depends on the form of the explanations, and in turn this shapes how to update the model too (discussed in Section 2.2.3). For text classification, most EBHD papers asked humans to decide which words (WO) in the explanation (considered important by the model) are in fact relevant or irrelevant (Kulesza et al., 2010; Ribeiro et al., 2016; Teso and Kersting, 2019). Some papers even allowed humans to adjust the word importance scores (WS) (Kulesza et al., 2009, 2015). This is analogous to specifying relevancy scores for example-based explanations (ES) in Zylberajch et al. (2021). Meanwhile, feedback at the level of learned features (FE) (i.e., the internal neurons in the model) and learned rules (RU) rather than individual words, was asked in Lertvittayakumjorn et al. (2020) and Ribeiro et al. (2018b), respectively. Additionally, humans may be asked to check the predicted labels (Kulesza et al., 2009; Smith-Renner et al., 2020) or even the ground truth labels (collectively noted as LB in Table 1) (Koh and Liang, 2017; Khanna et al., 2019; Han and Ghosh, 2020). Targeting the table question answering, Cho et al. (2019) asked humans to identify where in the table and

the question the model should focus (AT). This is analogous to identifying relevant words to attend for text classification.

It is likely that identifying important parts in the input is sufficient to make the model accomplish simple text classification tasks. However, this might not be enough for complex tasks that require reasoning. Recently, Yao et al. (2021) asked humans to provide, as feedback, compositional explanations to show how the humans would reason (RE) about the models' failure cases. An example of the feedback for a hate speech detection is "Because X is the word dumb, Y is a hateful word, and X is directly before Y , the attribution scores of both X and Y as well as the interaction score between X and Y should be increased". To acquire richer information like this as feedback, their framework requires more expertise from the feedback providers. In the future, it would be interesting to explore how we can collect and utilize other forms of feedback, for example, natural language feedback (Camburu et al., 2018), new training examples (Fiebrink et al., 2009), and other forms of decision rules used by humans (Carstens and Toni, 2017).

2.2.3 Updating the Model

Techniques to incorporate human feedback into the model can be categorized into three approaches.

(1) Directly adjust the model parameters (M). When the model is transparent and the explanation displays the model parameters in an intelligible way, humans can directly adjust the parameters based on their judgements. This idea was adopted by Kulesza et al. (2009, 2015) where humans can adjust a bar chart showing word importance scores, corresponding to the parameters of the underlying Naive Bayes model. In this special case, steps 2 and 3 in Figure 1 are combined into a single step. Besides, human feedback can be used to modify the model parameters indirectly. For example, Smith-Renner et al. (2020) increased a word weight in the Naive Bayes model by 20% for the class that the word supported, according to human feedback, and reduced the weight by 20% for the opposite class (binary classification). This choice gives good results, although it is not clear why and whether 20% is the best choice here.

Overall, this approach is fast because it does not require model retraining. However, it is important

to ensure that the adjustments made by humans generalize well to all examples. Therefore, the system should update the overall results (e.g., performance metrics, predictions, and explanations) in real time after applying any adjustment, so the humans can investigate the effects and further adjust the model parameters (or undo the adjustments) if necessary. This agrees with the correctness principles proposed by Kulesza et al. (2015) that the system should be actionable and reversible, honor user feedback, and show incremental changes.

(2) Improve the training data (D). We can use human feedback to improve the training data and retrain the model to fix bugs. This approach includes correcting mislabeled training examples (Koh and Liang, 2017; Han and Ghosh, 2020), assigning noisy labels to unlabeled examples (Yao et al., 2021), removing irrelevant words from input texts (Ribeiro et al., 2016), and creating augmented training examples to reduce the effects of the artifacts (Ribeiro et al., 2018b; Teso and Kersting, 2019; Zylberajch et al., 2021). As this approach modifies the training data only, it is applicable to any model regardless of the model complexity.

(3) Influence the training process (T). Another approach is to influence the (re-)training process in a way that the resulting model will behave as the feedback suggests. This approach could be either model-specific (such as attention supervision) or model-agnostic (such as user co-training). Cho et al. (2019) used human feedback to supervise attention weights of the model. Similarly, Yao et al. (2021) added a loss term to regularize explanations guided by human feedback. Stumpf et al. (2009) proposed (i) constraint optimization, translating human feedback into constraints governing the training process and (ii) user co-training, using feedback as another classifier working together with the main ML model in a semi-supervised learning setting. Lertvittayakumjorn et al. (2020) disabled some learned features deemed irrelevant, based on the feedback, and re-trained the model, forcing it to use only the remaining features. With many techniques available, however, there has not been a study testing which technique is more appropriate for which task, domain, or model architecture. The comparison issue is one of the

open problems for EBHD research (to be discussed in Section 4).

2.2.4 Iteration

The debugging workflow (explain, feedback, and update) can be done iteratively to gradually improve the model where the presented explanation changes after the model update. This allows humans to fix vital bugs first and finer bugs in later iterations, as reflected in Ribeiro et al. (2016) and Koh and Liang (2017) via the performance plots. However, the interactive process could be susceptible to *local decision pitfalls* where local improvements for individual predictions could add up to inferior overall performance (Wu et al., 2019). So, we need to ensure that the update in the current iteration is generally favorable and does not overwrite the good effects of previous updates.

2.3 Experimental Setting

To conduct experiments, some studies in Table 1 selected human participants (SP) to be their feedback providers. The selected participants could be people without ML/NLP knowledge (Kulesza et al., 2010, 2015) or with ML/NLP knowledge (Ribeiro et al., 2018b; Zylberajch et al., 2021) depending on the study objectives and the complexity of the feedback process. Early work even conducted experiments with the participants in-person (Stumpf et al., 2009; Kulesza et al., 2009, 2015). Although this limited the number of participants (to less than 100), the researchers could closely observe their behaviors and gain some insights concerning human-computer interaction.

By contrast, some used a crowdsourcing platform, Amazon Mechanical Turk⁵ in particular, to collect human feedback for debugging the models. Crowdsourcing (CS) enables researchers to conduct experiments at a large scale; however, the quality of human responses could be varying. So, it is important to ensure some quality control such as specifying required qualifications (Smith-Renner et al., 2020), using multiple annotations per question (Lertvittayakumjorn et al., 2020), having a training phase for participants, and setting up some obvious questions to check if the participants are paying attention to the tasks (Egelman et al., 2014).

⁵<https://www.mturk.com/>.

Finally, simulation (SM), without real humans involved but using oracles as human feedback instead, has also been considered (for the purpose of testing the EBHD framework only). For example, Teso and Kersting (2019) set 20% of input words as relevant using feature selection. These were used to respond to post-hoc explanations, that is, top k words selected by LIME. Koh and Liang (2017) simulated mislabeled examples by flipping the labels of a random 10% of the training data. So, when the explanation showed suspicious training examples, the true labels could be used to provide feedback. Compared to the other settings, simulation is faster and cheaper, yet its results may not reflect the effectiveness of the framework when deployed with real humans. Naturally, human feedback is sometimes inaccurate and noisy, and humans could also be interrupted or frustrated while providing feedback (Amershi et al., 2014). These factors, discussed in detail in the next Section, cannot be thoroughly studied in only simulated experiments.

3 Research on Human Factors

Though the major goal of EBHD is to improve models, we cannot disregard the effect on feedback providers of the debugging workflow. In this Section, we compile findings concerning how explanations and feedback could affect the humans, discussed along five dimensions: model understanding, willingness, trust, frustration, and expectation. Although some of the findings were not derived in NLP settings, we believe that they are generalizable and worth discussing in the context of EBHD.

3.1 Model Understanding

So far, we have used explanations as means to help humans understand models and conduct informed debugging. Hence, it is important to verify, at least preliminarily, that the explanations help feedback providers form an accurate understanding of how the models work. This is an important prerequisite towards successful debugging.

Existing studies have found that some explanation forms are more conducive to developing model understanding in humans than others. Stumpf et al. (2009) found that rule-based and keyword-based explanations were easier to understand than similarity-based explanations (i.e., explaining by similar examples in the training

data). Also, they found that some users did not understand why the absence of some words could make the model become more certain about its predictions. Lim et al. (2009) found that explaining why the system behaved and did not behave in a certain way resulted in good user understanding of the system, though the former way of explanation (why) was more effective than the latter (why not). Cheng et al. (2019) reported that interactive explanations could improve users' comprehension on the model better than static explanations, although the interactive way took more time. In addition, revealing inner workings of the model could further help understanding; however, it introduced additional cognitive workload that might make participants doubt whether they really understood the model well.

3.2 Willingness

We would like humans to provide feedback for improving models, but do humans naturally want to? Prior to the emerging of EBHD, studies found that humans are not willing to be constantly asked about labels of examples as if they were just simple oracles (Cakmak et al., 2010; Guillory and Bilmes, 2011). Rather, they want to provide more than just data labels after being given explanations (Amershi et al., 2014; Smith-Renner et al., 2020). By collecting free-form feedback from users, Stumpf et al. (2009) and Ghai et al. (2021) discovered various feedback types. The most prominent ones include removing-adding features (words), tuning weights, and leveraging feature combinations. Stumpf et al. (2009) further analyzed categories of background knowledge underlying the feedback and found, in their experiment, that it was mainly based on commonsense knowledge and English language knowledge. Such knowledge may not be efficiently injected into the model if we exploit human feedback that contains only labels. This agrees with some participants in Smith-Renner et al. (2020), who described their feedback as inadequate when they could only confirm or correct predicted labels.

Although human feedback beyond labels contains helpful information, it is naturally neither complete nor precise. Ghai et al. (2021) observed that human feedback usually focuses on a few features that are most different from human expectation, ignoring the others. Also, they found that humans, especially lay people, are not good at correcting model explanations quantitatively (e.g.,

adjusting weights). This is consistent with the findings of Miller (2019) that human explanations are selective (in a biased way) and rarely refer to probabilities but express causal relationships instead.

3.3 Trust

Trust (as well as frustration and expectation, discussed next) is an important issue when the system end users are feedback providers in the EBHD framework. It has been discussed widely that explanations engender human trust in AI systems (Pu and Chen, 2006; Lipton, 2018; Toreini et al., 2020). This trust may be misplaced at times. Showing more detailed explanations can cause users to over-rely on the system, leading to misuse where users agree with incorrect system predictions (Stumpf et al., 2016). Moreover, some users may over trust the explanations (without fully understanding them) only because the tools generating them are publicly available, widely used, and showing appealing visualizations (Kaur et al., 2020).

However, recent research reported that explanations do not necessarily increase trust and reliance. Cheng et al. (2019) found that, even though explanations help users comprehend systems, they cannot increase human trust in using the systems in high-stakes applications involving lots of qualitative factors, such as graduate school admissions. Smith-Renner et al. (2020) reported that explanations of low-quality models decrease trust and system acceptance as they reveal model weaknesses to the users. According to Schramowski et al. (2020), despite correct predictions, the trust still drops if the users see from the explanations that the model relies on the wrong reasons. These studies go along with a perspective by Zhang et al. (2020b) that explanations should help calibrate user perceptions to the model quality, signaling whether the users should trust or distrust the AI. Although, in some cases, explanations successfully warned users of faulty models (Ribeiro et al., 2016), this is not easy when the model flaws are not obvious (Zhang et al., 2020b; Lertvittayakumjorn and Toni, 2019).

Besides explanations, the effect of feedback on human trust is quite inconclusive according to some (but fewer) studies. On one hand, Smith-Renner et al. (2020) found that, after lay humans see explanations of low-quality models

and lose their trust, the ability to provide feedback makes human trust and acceptance rally, remedying the situation. In contrast, Honeycutt et al. (2020) reported that providing feedback decreases human trust in the system as well as their perception of system accuracy no matter whether the system truly improves after being updated or not.

3.4 Frustration

Working with explanations can cause frustration sometimes. Following the discussion on trust, explanations of poor models increase user frustration (as they reveal model flaws), whereas the ability to provide feedback reduces frustration. Hence, in general situations, the most frustrating condition is showing explanations to the users without allowing them to give feedback (Smith-Renner et al., 2020).

Another cause of frustration is the risk of detailed explanations overloading users (Narayanan et al., 2018). This is especially a crucial issue for inherently interpretable models where all the internal workings can be exposed to the users. Though presenting all the details is comprehensive and faithful, it could create barriers for lay users (Gershon, 1998). In fact, even ML experts may feel frustrated if they need to understand a decision tree with a depth of ten or more. Poursabzi-Sangdeh et al. (2018) found that showing all the model internals undermined users' ability to detect flaws in the model, likely due to information overload. So, they suggested that model internals should be revealed only when the users request to see them.

3.5 Expectation

Smith-Renner et al. (2020) observed that some participants expected the model to improve after the session where they interacted with the model, regardless of whether they saw explanations or gave feedback during the interaction session. EBHD should manage these expectations properly. For instance, the system should report changes or improvements to users after the model gets updated. It would be better if the changes can be seen incrementally in real time (Kulesza et al., 2015).

3.6 Summary

Based on the findings on human factors reviewed in this Section, we summarize suggestions for effective EBHD as follows.

Feedback Providers. Buggy models usually lead to implausible explanations, adversely affecting human trust in the system. Also, it is not yet clear whether giving feedback increases or decreases human trust. So, it is safer to let the developers or domain experts in the team (rather than end users) be the feedback providers. For some kinds of bugs, however, feedback from end users is essential for improving the model. To maintain their trust, we may collect their feedback implicitly (e.g., by inferring from their interactions with the system after showing them the explanations (Honeycutt et al., 2020)) or collect the feedback without telling them that the explanations are of the production system (e.g., by asking them to answer a separate survey). All in all, we need different strategies to collect feedback from different stakeholders.

Explanations. We should avoid using forms of explanations that are difficult to understand, such as similar training examples and absence of some keywords in inputs, unless the humans are already trained to interpret them. Also, too much information should be avoided as it could overload the humans; instead, humans should be allowed to request more information if they are interested, for example, by using interactive explanations (Dejl et al., 2021).

Feedback. Given that human feedback is not always complete, correct, or accurate, EBHD should use it with care, for example, by relying on collective feedback rather than individual feedback and allowing feedback providers to verify and modify their feedback before applying it to update the model.

Update. Humans, especially lay people, usually expect the model to improve over time after they give feedback. So, the system should display improvements after the model gets updated. Where possible, showing the changes incrementally in real time is preferred, as the feedback providers can check if their feedback works as expected or not.

4 Open Problems

This Section lists potential research directions and open problems for EBHD of NLP models.

4.1 Beyond English Text Classification

All papers in Table 1 conducted experiments only on English datasets. We acknowledge that qualitatively analyzing explanations and feedback in languages at which one is not fluent is not easy, not to mention recruiting human subjects who know the languages. However, we hope that, with more multilingual data publicly available (Wolf et al., 2020) and growing awareness in the NLP community (Bender, 2019), there will be more EBHD studies targeting other languages in the near future.

Also, most existing EBHD works target text classifiers. It would be interesting to conduct more EBHD work for other NLP tasks such as reading comprehension, question answering, and NLI, to see whether existing techniques still work effectively. Shifting to other tasks requires an understanding of specific bug characteristics in those tasks. For instance, unlike bugs in text classification, which are usually due to word artifacts, bugs in NLI concern syntactic heuristics between premises and hypotheses (McCoy et al., 2019). Thus, giving human feedback at word level may not be helpful, and more advanced methods may be needed.

4.2 Tackling More Challenging Bugs

Lakkaraju et al. (2020) remarked that the evaluation setup of existing EBHD work is often too easy or unrealistic. For example, bugs are obvious artifacts that could be removed using simple text pre-processing (e.g., removing punctuation and redacting named entities). Hence, it is not clear how powerful such EBHD frameworks are when dealing with real-world bugs. If bugs are not dominant and happen less often, global explanations may be too coarse-grained to capture them while many local explanations may be needed to spot a few appearances of the bugs, leading to inefficiency. As reported by Smith-Renner et al. (2020), feedback results in minor improvements when the model is already reasonably good.

Other open problems, whose solutions may help deal with challenging bugs, include the following. First, different people may give different feedback for the same explanation. As raised by Ghai et al. (2021), how can we integrate their feedback to obtain robust signals for model update? How should we deal with conflicts among feedback and training examples (Carstens and Toni, 2017)? Second, confirming or removing what the model has

learned is easier than injecting, into the model, new knowledge (which may not even be apparent in the explanations). How can we use human feedback to inject new knowledge, especially when the model is not transparent? Lastly, EBHD techniques have been proposed for tabular data and image data (Shao et al., 2020; Ghai et al., 2021; Popordanoska et al., 2020). Can we adapt or transfer them across modalities to deal with NLP tasks?

4.3 Analyzing and Enhancing Efficiency

Most selected studies focus on improving correctness of the model (e.g., by expecting a higher F1 or a lower bias after debugging). However, only some of them discuss efficiency of the proposed frameworks. In general, we can analyze the efficiency of an EBHD framework by looking at the efficiency of each main step in Figure 1. Step 1 generates the explanations, so its efficiency depends on the explanation method used and, in the case of local explanation methods, the number of local explanations needed. Step 2 lets humans give feedback, so its efficiency concerns the amount of time they spend to understand the explanations and to produce the feedback. Step 3 updates the model using the feedback, so its efficiency relates to the time used for processing the feedback and retraining the model (if needed). Existing work mainly reported efficiency of steps 1 or step 2. For instance, approaches using example-based explanations measured the improved performance with respect to the number of explanations computed (step 1) (Koh and Liang, 2017; Khanna et al., 2019; Han and Ghosh, 2020). Kulesza et al. (2015) compared the improved F1 of EBHD with the F1 of instance labeling given the same amount of time for humans to perform the task (step 2). Conversely, Yao et al. (2021) compared the time humans need to do EBHD versus instance labeling in order to achieve the equivalent degree of correctness improvement (step 2).

None of the selected studies considered the efficiency of the three steps altogether. In fact, the efficiency of steps 1 and 3 is important especially for black box models where the cost of post-hoc explanation generation and model retraining is not negligible. It is even more crucial for iterative or responsive EBHD. Thus, analyzing and enhancing efficiency of EBHD frameworks (for both machine and human sides) require further research.

4.4 Reliable Comparison Across Papers

User studies are naturally difficult to replicate as they are inevitably affected by choices of user interfaces, phrasing, population, incentives, and so forth (Lakkaraju et al., 2020). Further, research in ML rarely adopts practices from the human-computer interaction community (Abdul et al., 2018), limiting the possibility to compare across studies. Hence, most existing work only considers model performance before and after debugging or compares the results among several configurations of a single proposed framework. This leads to little knowledge about which explanation types or feedback mechanisms are more effective across several settings. Thus, one promising research direction would be proposing a standard setup or a benchmark for evaluating and comparing EBHD frameworks reliably across different settings.

4.5 Towards Deployment

So far, we have not seen EBHD research widely deployed in applications, probably due to its difficulty to set up the debugging aspects outside a research environment. One way to promote adoption of EBHD is to integrate EBHD frameworks into available visualization systems such as the Language Interpretability Tool (LIT) (Tenney et al., 2020), allowing users to provide feedback to the model after seeing explanations and supporting experimentation. Also, to move towards deployment, it is important to follow human-AI interaction guidelines (Amershi et al., 2019) and evaluate EBHD with potential end users, not just via simulation or crowdsourcing, since human factors play an important role in real situations (Amershi et al., 2014).

5 Conclusion

We presented a general framework of explanation-based human debugging (EBHD) of NLP models and analyzed existing work in relation to the components of this framework to illustrate the state-of-the-art in the field. Furthermore, we summarized findings on human factors with respect to EBHD, suggested design practices accordingly, and identified open problems for future studies. As EBHD is still an ongoing research topic, we hope that our survey will be helpful for guiding interested researchers and for examining future EBHD papers.

Acknowledgments

We would like to thank Marco Baroni (the Action Editor) and anonymous reviewers for very helpful comments. Also, we thank Brian Roark and Cindy Robinson for their technical support concerning the submission system. Additionally, the first author wishes to thank the support from Anandamahidol Foundation, Thailand.

References

- Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–18. <https://doi.org/10.1145/3173574.3174156>
- Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6:52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Julius Adebayo, Michael Muelly, Ilaria Liccardi, and Been Kim. 2020. Debugging tests for model explanations. In *Advances in Neural Information Processing Systems*.
- Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4):105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13. <https://doi.org/10.1145/3290605.3300233>
- Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2016. Explaining predictions of non-linear classifiers in NLP. In *Proceedings of the 1st*

- Workshop on Representation Learning for NLP*, pages 1–7, Berlin, Germany. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W16-1601>
- Yanzhe Bekkemoen and Helge Langseth. 2021. Correcting classification: A Bayesian framework using explanation feedback to improve classification abilities. *arXiv preprint arXiv:2105.02653*.
- Emily Bender. 2019. The #benderrule: On naming the languages we study and why it matters. *The Gradient*.
- Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José MF Moura, and Peter Eckersley. 2020. Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 648–657. <https://doi.org/10.1145/3351095.3375624>
- Gabriel Cadamuro, Ran Gilad-Bachrach, and Xiaojin Zhu. 2016. Debugging machine learning models. In *ICML Workshop on Reliable Machine Learning in the Wild*.
- Maya Cakmak, Crystal Chao, and Andrea L. Thomaz. 2010. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118. <https://doi.org/10.1109/TAMD.2010.2051030>
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. E-SNLI: Natural language inference with natural language explanations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 9560–9572.
- Lucas Carstens and Francesca Toni. 2017. Using argumentation to improve classification in natural language problems. *ACM Transactions on Internet Technology (TOIT)*, 17(3):1–23. <https://doi.org/10.1145/3017679>
- Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. 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*, pages 1721–1730. <https://doi.org/10.1145/2783258.2788613>
- Gromit Yeuk-Yin Chan, Jun Yuan, Kyle Overton, Brian Barr, Kim Rees, Luis Gustavo Nonato, Enrico Bertini, and Claudio T. Silva. 2020. Subplex: Towards a better understanding of black box model explanations at the subpopulation level. *arXiv preprint arXiv:2007.10609*.
- Hao-Fei Cheng, Ruotong Wang, Zheng Zhang, Fiona O’Connell, Terrance Gray, F. Maxwell Harper, and Haiyi Zhu. 2019. Explaining decision-making algorithms through UI: Strategies to help non-expert stakeholders. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12. <https://doi.org/10.1145/3290605.3300789>
- Minseok Cho, Reinald Kim Amplayo, Seung-won Hwang, and Jonghyuck Park. 2018. Adversarial tableqa: Attention supervision for question answering on tables. In *Proceedings of The 10th Asian Conference on Machine Learning*, volume 95 of *Proceedings of Machine Learning Research*, pages 391–406. PMLR.
- Minseok Cho, Gyeongbok Lee, and Seung-won Hwang. 2019. Explanatory and actionable debugging for machine learning: A tableqa demonstration. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1333–1336.
- Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction*, 18(5):455. <https://doi.org/10.1007/s11257-008-9051-3>
- 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*, pages 447–459, Suzhou, China. Association for Computational Linguistics.
- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* '19*, pages 120–128, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287572>
- Adam Dejl, Peter He, Pranav Mangal, Hasan Mohsin, Bogdan Surdu, Eduard Voinea, Emanuele Albini, Piyawat Lertvittayakumjorn, Antonio Rago, and Francesca Toni. 2021. Argflow: A toolkit for deep argumentative explanations for neural networks. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1761–1763.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Serge Egelman, Ed H. Chi, and Steven Dow. 2014. Crowdsourcing in HCI research. In *Ways of Knowing in HCI*, Springer, pages 267–289.
- Rebecca Fiebrink, Dan Trueman, and Perry R. Cook. 2009. A metainstrument for interactive, on-the-fly machine learning. In *Proceedings of NIME*.
- Nahum Gershon. 1998. Visualization of an imperfect world. *IEEE Computer Graphics and Applications*, 18(4):43–45. <https://doi.org/10.1109/38.689662>
- Bhavya Ghai, Q. Vera Liao, Yunfeng Zhang, Rachel Bellamy, and Klaus Mueller. 2021. Explainable active learning (xal) toward ai explanations as interfaces for machine teachers. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–28. <https://doi.org/10.1145/3432934>
- Filip Graliński, Anna Wróblewska, Tomasz Stanisławek, Kamil Grabowski, and Tomasz Górecki. 2019. GEval: Tool for debugging NLP datasets and models. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 254–262, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-4826>
- Andrew Guillory and Jeff Bilmes. 2011. Simultaneous learning and covering with adversarial noise. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, pages 369–376, Madison, WI, USA. Omnipress.
- Han Guo, Nazneen Fatema Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. 2020. Fastif: Scalable influence functions for efficient model interpretation and debugging. *arXiv preprint arXiv:2012.15781*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-2017>
- Xiaochuang Han, Byron C. Wallace, and Yulia Tsvetkov. 2020. Explaining black box predictions and unveiling data artifacts through influence functions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5553–5563, Online. Association for Computational Linguistics.
- Xing Han and Joydeep Ghosh. 2020. Model-agnostic explanations using minimal forcing subsets. *arXiv preprint arXiv:2011.00639*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*,

- 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>, PubMed: 9377276
- Donald Honeycutt, Mahsan Nourani, and Eric Ragan. 2020. Soliciting human-in-the-loop user feedback for interactive machine learning reduces user trust and impressions of model accuracy. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 63–72.
- Benjamin Hoover, Hendrik Strobelt, and Sebastian Gehrmann. 2020. exBERT: A visual analysis tool to explore learned representations in transformer models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 187–196, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-demos.22>
- Maximilian Idahl, Lijun Lyu, Ujwal Gadiraju, and Avishek Anand. 2021. Towards benchmarking the utility of explanations for model debugging. *arXiv preprint arXiv:2105.04505*. <https://doi.org/10.18653/v1/2021.trustnlp-1.8>
- Alon Jacovi, Ana Marasović, Tim Miller, and Yoav Goldberg. 2020. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. *arXiv preprint arXiv:2010.07487*. <https://doi.org/10.1145/3442188.3445923>
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D17-1215>
- Xisen Jin, Zhongyu Wei, Junyi Du, Xiangyang Xue, and Xiang Ren. 2020. Towards hierarchical importance attribution: Explaining compositional semantics for neural sequence models. In *International Conference on Learning Representations*.
- David Johnson, Giuseppe Carenini, and Gabriel Murray. 2020. Njm-vis: Interpreting neural joint models in NLP. In *Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI '20*, pages 28–296, Association for Computing Machinery. New York, NY, USA. <https://doi.org/10.1145/3377325.3377513>
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics. <https://doi.org/10.18653/v1/E17-2068>
- Patrick Juola. 2007. Future trends in authorship attribution. In *IFIP International Conference on Digital Forensics*, pages 119–132. Springer. https://doi.org/10.1007/978-0-387-73742-3_8
- Daniel Kang, Deepti Raghavan, Peter Bailis, and Matei Zaharia. 2018. Model assertions for debugging machine learning. In *NeurIPS ML Sys Workshop*.
- Harmanpreet Kaur, Harsha Nori, Samuel Jenkins, Rich Caruana, Hanna Wallach, and Jennifer Wortman Vaughan. 2020. Interpreting interpretability: Understanding data scientists’ use of interpretability tools for machine learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Rajiv Khanna, Been Kim, Joydeep Ghosh, and Sanmi Koyejo. 2019. Interpreting black box predictions using fisher kernels. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3382–3390. PMLR.
- Sung Wook Kim, Iljeok Kim, Jonghwan Lee, and Seungchul Lee. 2021. Knowledge integration into deep learning in dynamical systems: An overview and taxonomy. *Journal of Mechanical Science and Technology*, pages 1–12.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*,

- pages 1746–1751, Doha, Qatar. Association for Computational Linguistics. <https://doi.org/10.3115/v1/D14-1181>
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pages 1885–1894. PMLR.
- Josua Krause, Adam Perer, and Kenney Ng. 2016. Interacting with predictions: Visual inspection of black-box machine learning models. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5686–5697. <https://doi.org/10.1145/2858036.2858529>
- Sanjay Krishnan and Eugene Wu. 2017. Palm: Machine learning explanations for iterative debugging. In *Proceedings of the 2nd Workshop on Human-In-the-Loop Data Analytics*, pages 1–6. <https://doi.org/10.1145/3077257.3077271>
- Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*, pages 126–137. <https://doi.org/10.1145/2678025.2701399>
- Todd Kulesza, Simone Stumpf, Margaret Burnett, Weng-Keen Wong, Yann Riche, Travis Moore, Ian Oberst, Amber Shinsel, and Kevin McIntosh. 2010. Explanatory debugging: Supporting end-user debugging of machine-learned programs. In *2010 IEEE Symposium on Visual Languages and Human-Centric Computing*, pages 41–48. IEEE. <https://doi.org/10.1109/VLHCC.2010.15>
- Todd Kulesza, Weng-Keen Wong, Simone Stumpf, Stephen Perona, Rachel White, Margaret M. Burnett, Ian Oberst, and Andrew J. Ko. 2009. Fixing the program my computer learned: Barriers for end users, challenges for the machine. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, pages 187–196. <https://doi.org/10.1145/1502650.1502678>
- Vivian Lai, Han Liu, and Chenhao Tan. 2020. “why is’ chicago’deceptive?” towards building model-driven tutorials for humans. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13.
- Vivian Lai and Chenhao Tan. 2019. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 29–38.
- Himabindu Lakkaraju, Julius Adebayo, and Sameer Singh. 2020. Explaining machine learning predictions: State-of-the-art, challenges, and opportunities. NeurIPS 2020 Tutorial.
- Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 331–339. <https://doi.org/10.1016/B978-1-55860-377-6.50048-7>
- Piyawat Lertvittayakumjorn, Ivan Petej, Yang Gao, Yamuna Krishnamurthy, Anna Van Der Gaag, Robert Jago, and Kostas Stathis. 2021. Supporting complaints investigation for nursing and midwifery regulatory agencies. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 81–91, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.acl-demo.10>
- Piyawat Lertvittayakumjorn, Lucia Specia, and Francesca Toni. 2020. FIND: Human-in-the-loop debugging deep text classifiers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 332–348, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.24>
- Piyawat Lertvittayakumjorn and Francesca Toni. 2019. Human-grounded evaluations of explanation methods for text classification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*,

- pages 5195–5205, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1523>
- Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2119–2128.
- Zachary C. Lipton. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57. <https://doi.org/10.1145/3236386.3241340>
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Raoni Lourenço, Juliana Freire, and Dennis Shasha. 2020. Bugdoc: A system for debugging computational pipelines. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 2733–2736. <https://doi.org/10.1145/3318464.3384692>
- Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30:4765–4774.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1334>
- Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267:1–38. <https://doi.org/10.18653/v1/P19-1334>
- Yao Ming, Panpan Xu, Huamin Qu, and Liu Ren. 2019. Interpretable and steerable sequence learning via prototypes. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 903–913. <https://doi.org/10.1145/3292500.3330908>
- Brad A. Myers, David A. Weitzman, Andrew J. Ko, and Duen H. Chau. 2006. Answering why and why not questions in user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 397–406. <https://doi.org/10.1145/1124772.1124832>
- Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finale Doshi-Velez. 2018. How do humans understand explanations from machine learning systems? An evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682*.
- Devi Parikh and C. Zitnick. 2011. Human-debugging of machines. *NIPS WCSSWC*, 2(7):3. <https://doi.org/10.18653/v1/D18-1302>
- Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. Reducing gender bias in abusive language detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2799–2804, Brussels, Belgium. Association for Computational Linguistics.
- Teodora Popordanoska, Mohit Kumar, and Stefano Teso. 2020. Machine guides, human supervises: Interactive learning with global explanations. *arXiv preprint arXiv:2009.09723*.
- Forough Poursabzi-Sangdeh, Daniel G. Goldstein, Jake M. Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. 2018. Manipulating and measuring model interpretability. *arXiv preprint arXiv:1802.07810*.
- Pearl Pu and Li Chen. 2006. Trust building with explanation interfaces. In *Proceedings of the 11th international conference on Intelligent user interfaces*, pages 93–100.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, pages 1–26.

- Marco Tulio Ribeiro, Carlos Guestrin, and Sameer Singh. 2019. Are red roses red? Evaluating consistency of question-answering models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6174–6184, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1621>
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018a. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018b. Semantically equivalent adversarial rules for debugging NLP models. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1079>
- 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*, pages 4902–4912, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.442>
- Laura Rieger, Chandan Singh, William Murdoch, and Bin Yu. 2020. Interpretations are useful: Penalizing explanations to align neural networks with prior knowledge. In *International Conference on Machine Learning*, pages 8116–8126. PMLR.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A primer in bertology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866. https://doi.org/10.1162/tacl_a_00349
- Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Laura von Rueden, Sebastian Mayer, Katharina Beckh, Bogdan Georgiev, Sven Giesselbach, Raoul Heese, Birgit Kirsch, Michal Walczak, Julius Pfrommer, Annika Pick, Rajkumar Ramamurthy, Michal Walczak, Jochen Garcke, Christian Bauckhage, and Jannis Schuecker. 2021. Informed machine learning—a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Transactions on Knowledge and Data Engineering*. <https://doi.org/10.1109/TKDE.2021.3079836>
- Eldon Schoop, Forrest Huang, and Björn Hartmann. 2020. Scram: Simple checks for realtime analysis of model training for non-expert ml programmers. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–10. <https://doi.org/10.1145/3334480.3382879>
- Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. 2020. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nature Machine Intelligence*, 2(8):476–486. <https://doi.org/10.1038/s42256-020-0212-3>
- Daniel Selsam, Percy Liang, and David L. Dill. 2017. Developing bug-free machine learning systems with formal mathematics. In *International Conference on Machine Learning*, pages 3047–3056. PMLR.
- Xiaoting Shao, Tjitze Rienstra, Matthias Thimm, and Kristian Kersting. 2020. Towards understanding and arguing with classifiers: Recent progress. *Datenbank-Spektrum*, 20(2):171–180. <https://doi.org/10.1007/s13222-020-00351-x>

- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. 2017. Smoothgrad: Removing noise by adding noise. *arXiv preprint arXiv:1706.03825*.
- Alison Smith-Renner, Ron Fan, Melissa Birchfield, Tongshuang Wu, Jordan Boyd-Graber, Daniel S. Weld, and Leah Findlater. 2020. No explainability without accountability: An empirical study of explanations and feedback in interactive ml. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13. <https://doi.org/10.1145/3313831.3376624>
- Simone Stumpf, Adrian Bussone, and Dymrna O’Sullivan. 2016. Explanations considered harmful? User interactions with machine learning systems. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*.
- Simone Stumpf, Vidya Rajaram, Lida Li, Weng-Keen Wong, Margaret Burnett, Thomas Dietterich, Erin Sullivan, and Jonathan Herlocker. 2009. Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies*, 67(8):639–662. <https://doi.org/10.1016/j.ijhcs.2009.03.004>
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Ian Tenney, James Wexler, Jasmijn Bastings, Tolga Bolukbasi, Andy Coenen, Sebastian Gehrmann, Ellen Jiang, Mahima Pushkarna, Carey Radebaugh, Emily Reif, and Ann Yuan. 2020. The language interpretability tool: Extensible, interactive visualizations and analysis for NLP models. <https://doi.org/10.18653/v1/2020.emnlp-demos.15>
- Stefano Teso and Kristian Kersting. 2019. Explanatory interactive machine learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 239–245. <https://doi.org/10.1145/3306618.3314293>
- Ehsan Toreini, Mhairi Aitken, Kovila Coopamootoo, Karen Elliott, Carlos Gonzalez Zelaya, and Aad van Moorsel. 2020. The relationship between trust in AI and trustworthy machine learning technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 272–283. <https://doi.org/10.1145/3351095.3372834>
- Zijie J. Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. 2021. Putting humans in the natural language processing loop: A survey. *arXiv preprint arXiv:2103.04044*.
- Thomas Wolf, Quentin Lhoest, Patrick von Platen, Yacine Jernite, Mariama Drame, Julien Plu, Julien Chaumond, Clement Delangue, Clara Ma, Abhishek Thakur, Suraj Patil, Joe Davison, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angie McMillan-Major, Simon Brandeis, Sylvain Gugger, François Lagunas, Lysandre Debut, Morgan Funtowicz, Anthony Moi, Sasha Rush, Philipp Schmid, Pierric Cistac, Victor Muštar, Jeff Boudier, and Anna Tordjmann. 2020. Datasets. *GitHub*. Note: <https://github.com/huggingface/datasets,1>.
- Tongshuang Wu, Daniel S. Weld, and Jeffrey Heer. 2019. Local decision pitfalls in interactive machine learning: An investigation into feature selection in sentiment analysis. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(4):1–27. <https://doi.org/10.1145/3319616>
- Yao, Huihan, Chen, Ying, Ye, Qinyuan, Jin, Xisen, and Ren, Xiang. 2021. Refining Language Models with Compositional Explanations. *Advances in Neural Information Processing Systems*, 34.
- Roosbeh Yousefzadeh and Dianne P. O’Leary. 2019. Debugging trained machine learning models using flip points. In *ICLR 2019 Debugging Machine Learning Models Workshop*.
- Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2020a. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41. <https://doi.org/10.1145/3374217>

- Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020b. Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 295–305. <https://doi.org/10.1145/3351095.3372852>
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4995–5004. <https://doi.org/10.1109/CVPR.2016.540>
- Hugo Zylberajch, Piyawat Lertvittayakumjorn, and Francesca Toni. 2021. HILDIF: Interactive debugging of NLI models using influence functions. In *Proceedings of the First Workshop on Interactive Learning for Natural Language Processing*, pages 1–6, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.internlp-1.1>