

Towards Human-Centred Explainability Benchmarks For Text Classification

Viktor Schlegel, Erick Mendez-Guzman and Riza Batista-Navarro

Department of Computer Science, University of Manchester
Manchester, United Kingdom
{viktor.schlegel, erick.mendezguzman, riza.batista}@manchester.ac.uk

Abstract

Progress on many Natural Language Processing (NLP) tasks, such as text classification, is driven by objective, reproducible and scalable evaluation via publicly available benchmarks. However, these are not always representative of real-world scenarios where text classifiers are employed, such as sentiment analysis or misinformation detection.

In this position paper, we put forward two points that aim to alleviate this problem. First, we propose to extend text classification benchmarks to evaluate the *explainability* of text classifiers. We review challenges associated with objectively evaluating the capabilities to produce valid explanations which leads us to the second main point: We propose to ground these benchmarks in human-centred applications, for example by using social media, gamification or to learn explainability metrics from human judgements.

Introduction

Recent advancements in deep learning research have yielded impressive results for many Natural Language Processing (NLP) tasks, with text classification (Kowsari et al. 2019; Otter, Medina, and Kalita 2020) being no exception. A popular application domain for text classification systems is social media, as it features many high-stake use cases (Minaee et al. 2021). Deep neural networks have been used for suicidal ideation detection and depression identification (Sawhney et al. 2018; Tadesse et al. 2019), analysing public perceptions of physical distancing interventions during the COVID-19 outbreak (Raamkumar et al. 2020), and identification of modern slavery (Guzman, Schlegel, and Batista-Navarro 2022).

Their impressive performance can be attributed to two factors: on the one hand, conceptually simple general-purpose neural architectures that provide expressive contextualised embeddings optimised on large corpora (Devlin et al. 2019) are readily available to be fine-tuned to a target application domain. On the other hand, the existence of these standardised approaches facilitates the development of benchmarks (Thorne et al. 2018; Rajpurkar et al. 2016; Wang et al. 2018, 2019), that allow us to perform scalable, automated and reproducible evaluation. This, in turn, fosters healthy competition, driven by advantages associated with developing approaches that perform best on these benchmarks, such as increased visibility in the research com-

munity (Linzen 2020; Church and Hestness 2019). In this setting, progress can naturally only be measured (and thus achieved) towards those objectives that are explicitly modelled with corresponding performance metrics as the target of evaluation. This development can be summarised with the proverb “you reap what you sow”: providing the research community with a competitive and accessible benchmark that evaluates some metric A will result in approaches that are explicitly optimised for A , gradually pushing A ’s state-of-the-art measure.

The vast majority of popular benchmarks focus on intrinsic task performance only, reflected in metrics such as accuracy (Ribeiro et al. 2020). This may explain more recent findings which suggest that evaluated approaches tend to overfit to these benchmarks by relying on spurious correlations in data in favour of reaching a high score, rather than successfully solving the represented task (Gururangan et al. 2018). This is evidenced by their moderate generalisation performance to examples drawn outside of a benchmark’s distribution (Schlegel, Nenadic, and Batista-Navarro 2020), or their brittleness towards adversarial inputs (Wallace et al. 2019a). These objectives are of vital importance in real-world scenarios, however. For example, utilising a static dataset to evaluate a tool that detects misinformation on social media contradicts the fact that language use in social media changes over time (Eisenstein et al. 2014), and contemporary events may lead to claims regarding completely new topics. Similarly, an automated tool should be robust to perturbations in the lexical form of how misinformation is conveyed.

Different solutions have been proposed to address these issues, such as “training set-free benchmarks” to prevent overfitting to benchmark data (Linzen 2020), or adversarial challenge sets that evaluate the robustness of classifiers to out-of-distribution examples (McCoy, Pavlick, and Linzen 2019). We propose to take a complementary route and to focus on the capability of classifiers to explain their predictions. In the context of the present position paper, explainability refers to the process of revealing the inner workings behind a model prediction for a single input text (Gunning et al. 2019; Danilevsky et al. 2020). In contrast, interpretability refers to a characteristic of a model that makes its global logic easily understandable by end-users (Atanasova et al. 2020). We discuss why explanations are important,

what makes explanations “good” and how their quality is evaluated and conclude with a set of specific research questions for combining human-grounded evaluation and large-scale objective evaluation in order to make these benchmarks reflective of real-world application scenarios.

Why are Explanations Important?

The first, perhaps more straightforward use case for explanations is their capability to contextualise predictions. Explanations are necessary for the current generation of state-of-the-art solutions to text classification problems, because unlike other algorithms, such as those based on symbolic rules, e.g. decision trees, or human-interpretable features, it is not directly possible to predict the behaviour of a neural network-based classifier (the *what*) or to explain it (the *why*) by looking at its performed computations (the *how*). Neural networks are agnostic to input data (e.g. the same algorithm can be used to classify text or detect cats in images); the properties of the resulting models are largely determined by the data they are optimised on. This is not necessarily a problem: if there is access to data that is fully representative of a well-defined task, as is the case with synthetic datasets (Weston et al. 2015), and a model that solves this task perfectly, one might argue that understanding its behaviour is secondary given the reliability of its performance.

As previously argued, this assumption appears to not hold for many text classification tasks due to the dynamic and compositional nature of language, and can be seen by the lack of out-of-distribution generalisation. In this scenario, an explanation can provide end-users the necessary context to trust or challenge the prediction (Glass, McGuinness, and Wolverson 2008). For example, a nonsensical explanation to a social media post classified as misinformation could hint at the fact that the prediction is wrong. In this context, explanations serve the same purpose as many other tools aimed at providing additional context, such as with model cards (Mitchell et al. 2019), behavioural testing suites (Ribeiro et al. 2020) or interpretability methods (Belinkov and Glass 2019). Different from those, explanations are arguably the easiest to understand for end users as they do not require any technical knowledge.

Beyond their use as a diagnostic tool, explanations can help to evaluate the classifier itself. This is intuitive: being able to explain something is a sufficient but not necessary indicator to understand it—to explain something you need to understand it first and good performance on the harder task of producing a good explanation implies robust performance on the classification task. This intuition is supported by the fact that jointly learning to both predict and explain, results in more robust classifiers that are less susceptible to weaknesses such as reliance on spurious correlations (Rajagopal et al. 2020; Dua, Singh, and Gardner 2020; Atanasova et al. 2020; Lei, Barzilay, and Jaakkola 2016).

Therefore, objective and large-scale evaluation of classifiers’ capabilities to produce explanations is desirable, and benchmarks that require approaches to explicitly optimise these capabilities, can contribute to the development of robust classifiers suitable for real-world applications. In the

following sections we discuss why such benchmarks have yet to gain traction.

What Makes a Good Explanation?

Even in practical terms, what constitutes a valid explanation is an open question. On a high level, different desiderata for evaluating explanations have been proposed: *plausibility*, or human-likeness describes whether humans are likely to be convinced by a machine-generated explanation. Put simply, an explanation is plausible if a majority of humans exposed to the explanation agree that it explains the prediction well. Furthermore, *faithfulness*, also referred to as fidelity, describes the extent to which explanations inform model predictions (Ross, Hughes, and Doshi-Velez 2017), which is important, because a classifier’s plausible explanation should in fact be reflected in its prediction.

Different approaches have been proposed for explaining the outcome of neural text classifiers. Explanations by example exemplify a classification by presenting most similar examples in the training data (Han, Wallace, and Tsvetkov 2020). Counter-factual explanations, rooted in the causal reasoning framework (Hendricks et al. 2018; Garg et al. 2019) identify which parts of the input need to be changed in order to arrive at a different explanation. Scientific explanations aim to link the input to the predicted label via reasoning chains and background knowledge (Jansen et al. 2018). Finally, *rationales* are text snippets that justify or warrant the predicted class.

In its simplest form, these are based on building connections between the input text and output labels and quantifying how much each element (words or n -grams) contributes to the final prediction (Atanasova et al. 2020). In the context of deep learning-based methods for text classification, these extractive post-hoc explanation methods commonly take the form of token-level importance scores (Danilevsky et al. 2020). While *soft* scoring models assign soft weights to token representations, and so one can extract highly weighted tokens as explanations, *hard* selection mechanisms discretely extract snippets from the input text as explanations (Atanasova et al. 2020; Danilevsky et al. 2020). Relaxing the requirement that rationales must be part of the input, free-form rationales take the form of short texts aimed at justifying the label (Liu, Yin, and Wang 2018). These can be generated with Natural Language Generation approaches, such as Sequence-to-Sequence models (Costa et al. 2018; Rajani et al. 2019).

How is Explainability Evaluated?

Concerning evaluation, there is little consensus on how to assess the quality of these explanations either for deployment or benchmarking (Doshi-Velez and Kim 2017). Current evaluation approaches typically fall into two categories. The first uses some formal definition of explainability desiderata as a quantifiable proxy for explanation quality and requires no human experiments (Nguyen 2018; Carton, Rathore, and Tan 2020). The second evaluates explainability in the context of an application using human experiments either for the practical application or a simplified version

of it (Lertvittayakumjorn and Toni 2019; Hase and Bansal 2020).

Proxy-based Evaluation Model-generated explanations have typically been evaluated by automatic proxy measures of similarity with human explanations (Hase and Bansal 2020). In this context, the term “proxy” denotes that an objective metric is used to—perhaps imperfectly—represent a more abstract concept, such as plausibility. For the discrete case, researchers have measured agreement using phrase-matching metrics such as BLEU and MAXSIM (Camburu et al. 2018; Clinciu, Eshghi, and Hastie 2021), Intersection-over-Union (IOU) on a token level to derive precision, recall and F1-scores (DeYoung et al. 2019). Meanwhile, for soft token scoring models, some studies have used the Area Under the Precision-Recall curve (AUPRC) constructed by sweeping a threshold over token scores (Samek et al. 2016).

While these proxy metrics are suitable for automated and scalable evaluation, there is no guarantee that they represent human intuition well. At their core, these metrics are evaluating lexical overlap and might give unreasonably high scores when predicted rationales are lexically similar yet semantically different or unreasonably low scores in the converse case. For text generation tasks, this results in poor correlation between human ratings and scores under these metrics (Wiseman, Shieber, and Rush 2017; Ma et al. 2019). Additionally, for generated explanations, trivial baselines, such as input copying, may outperform bespoke approaches (Krishna, Roy, and Iyyer 2021).

Furthermore, even though the evaluation metrics cited above provide a good starting point, they do not reveal model behaviour as they evaluate plausibility but not faithfulness (Hase et al. 2020; Carton, Rathore, and Tan 2020), because they do not reveal whether the produced explanations actually informs the classifier’s prediction. Despite the fact that the best way to measure explanation faithfulness for text classification remains an open question, most researchers investigating it have defined that a faithful explanation should have high sufficiency and comprehensiveness (DeYoung et al. 2019). Sufficiency measures how well explanations can provide nearly the same prediction as the whole input text, while comprehensiveness assesses how well machine-generated explanations include all relevant tokens (Liu, Yin, and Wang 2018).

More recently, literature presenting contradictory findings of using proxy-based metrics for evaluating machine-generated explanations, has emerged (Carton, Rathore, and Tan 2020; Lertvittayakumjorn and Toni 2019). Previous research has established that human explanations do not necessarily have high sufficiency and comprehensiveness, suggesting they may not be treated as ground truth baselines (Carton, Rathore, and Tan 2020).

The issue has grown in importance in light of recent evidence that indicates that automated metrics are highly model-dependent, which makes it challenging to compare explanations from different text classification models (Clinciu, Eshghi, and Hastie 2021). Moreover, metrics for assessing explanation faithfulness are unstable compared to model performance. For instance, sufficiency and compre-

hensiveness may continue to fluctuate significantly even when model accuracy has stabilised (Carton, Rathore, and Tan 2020).

Human-grounded Evaluation With a concrete application in mind, the best way to assess model explanations is to evaluate them for the specific task. If such evaluation with end-users is not feasible, simpler human-subject experiments may be a good alternative if they maintain the essence of the target application.

To date, several studies have begun to examine the use of human-grounded evaluation methods to assess explanation quality for NLP applications (Mohseni, Block, and Ragan 2018; Nguyen 2018; Lertvittayakumjorn and Toni 2019; Mohseni, Block, and Ragan 2021). Together, these studies indicate that human experiments can help to evaluate whether explanations can expose irrational behaviour of a text classifier, e.g. the exploitation of dataset-specific spurious correlations (Lertvittayakumjorn and Toni 2019), assess whether explanations are faithful concerning the predicted class (Mohseni, Block, and Ragan 2018, 2021) or to investigate uncertain predictions (Lertvittayakumjorn and Toni 2019). Human-centred experiments can also help to measure the users’ understanding of explanations, establish the usefulness of explanations for a specific task (Paleja et al. 2021), and assess user-reported trust as a proxy for explanation goodness (Lertvittayakumjorn and Toni 2019; Lage et al. 2019; Kocielnik, Amershi, and Bennett 2019).

How to Benchmark Explainability

We have reviewed the challenges of evaluating text classification systems robustly, objectively and at scale. Objective evaluation is exacerbated by spurious correlations in static benchmarks and changing requirements between training data and real-world application scenarios. To detect and alleviate these issues, we propose to expand text classification benchmarks by a harder task, **explanation reconstruction**. This forms a sufficient requirement for successful text classification, while doing so “for the right reasons”, rather than relying on superficial heuristics and shortcuts. Specifically, we discuss the tasks of rationale extraction and generation as conceptually simple yet flexible candidates. In this section, we outline the necessary steps to transition from performance-focussed benchmark suites to those that explicitly evaluate explainability.

In the **short-term**, we argue that rationales need to be incorporated into large-scale classification benchmarks. Conceptually, both generating and extracting rationales are well-explored supervised learning tasks (natural language generation and span extraction, respectively), so any classification benchmarks and classifiers they evaluate can be feasibly enriched by these tasks. *For new benchmarks*, required training and evaluation data can be collected by crowd-workers during annotation for the main task due to the conceptual simplicity of rationales, as they do not require additional knowledge, such as a training set to draw analogies from or additional resources to form reasoning chains. In other words, providing a justification does not significantly increase the cognitive load of annotators, because they already perform

this task when choosing the correct label, albeit implicitly (Zaidan, Eisner, and Piatko 2007). For existing benchmarks, rationale annotations can be gathered post-publication (Dua et al. 2019).

Evaluation of explainability in such benchmarks can be carried out with the existing proxy-based metrics based on overlap with human-annotated rationales.

In the **longer term**, however, proxy-based metrics only reflect general desiderata of explainability (Lertvitayakumjorn and Toni 2019), and are not aimed at establishing the *usefulness* of produced explanations. By usefulness we mean the level of utility an explanation brings in addition to the raw prediction, when supporting decision-making. For example, a human might report an explanation to a misleading claim to be plausible and faithful. However, this is not necessarily reflective of whether the same human would refrain from sharing the claim on social media. Thus, to better reflect real-world human-centred applications in benchmarks, requirements for valid explanations for different tasks and for different user groups need to be gathered. As argued previously, human-centred experiments are best to establish and represent these requirements. Unfortunately, these evaluation methods do not exhibit any of the features that make benchmarks popular: Due to the human involvement they are not automated, and they might not be easily scalable and reproducible due to the effort required. Every newly developed approach will require a costly re-run of the human-centred evaluation.

With the need of objective explanation evaluation and in light of the challenges associated with existing proxy-based approaches, we propose to develop evaluation metrics based on task-specific signals provided by humans, effectively simulating expensive human-centred experiments. We propose three possible ways to achieve and highlight associated practical questions that need to be considered.

Gamification: Gamification (“Beat the AI”) or dynamic benchmarks have been proposed more recently to combat the reliance on dataset-specific artefacts in static benchmarks (Nie et al. 2020; Kiela et al. 2021; Wallace et al. 2019b). In these settings, examples would only be accepted if they can “fool” some pre-trained classifier to produce the wrong label, while the majority of humans would still annotate the correct one. The game-like nature of the annotation task provides additional motivation, at least in the short term (Hamari, Koivisto, and Sarsa 2014). Similarly, this annotation paradigm can be extended to continuously evaluate explanations generated by optimised models, for example by accepting only those challenging examples, where an explanation generated by a system would be voted as low-quality by majority of human annotators, effectively serving as a human-in-the-loop explanation benchmark.

Learned explainability evaluation metrics: More recently, overlap-based metrics have been complemented by embedding-based metrics in an effort to improve the correlation between automatically measured and human-judged quality of generated text (Zhang et al. 2019; Sellam, Das, and Parikh 2020). Similar metrics can be developed to directly learn plausibility, faithfulness or usefulness of explanations based on human judgements, effectively approxi-

mately replicating human-centred experiments. Predicting these scores directly given an explanation, has the further advantage that explanations can potentially be evaluated *without gold standard reference explanations*. However, it is worth investigating, whether the learned explainability metrics generalise beyond the training/evaluation data or whether the metrics themselves are susceptible to spurious correlations. Finally, overlap-based metrics have the advantage of being interpretable. It is unclear whether this property is desired from learned metrics.

Use of Social Media: Finally, the dynamics of social media can be helpful to approximate how *useful* or convincing generated explanations are. Zellers et al. (2020) post machine-generated explanations to the REDDITADVICE forum and approximate their usefulness by the number of “up-votes” received. Similarly, researchers can expose generated explanations to the public by making use of other fora that have voting, liking or sharing mechanisms. For example, the utility of fact verification tools can be approximated by automatically identifying misleading tweets and commenting with generated explanations why they are misleading. Different explanation methods can then be compared by the number of retweets they can accumulate or by how much they manage to dampen the expected spread of the original claim. Studies like these will need to consider the trade-off between impartiality and potential harm. On the one hand prefacing generated explanations with the fact that they are machine-generated can have an impact on evaluation, while on the other hand, “incorrect explanations” could cause harm.

We hope this position paper will contribute to the ongoing dialogue around incorporating human-centred, application-grounded scalable evaluation of explainability capabilities into predominantly performance-based benchmarking of text classification approaches. This will contribute to the development of robust benchmarks that are not easily exploitable firstly, and secondly, are more representative of humans’ needs in real-world scenarios.

Acknowledgements

We thank the reviewers for their feedback. This work was partially funded by the European Union’s Horizon 2020 research and innovation action programme, via the AI4Media Open Call #1 issued and executed under the AI4Media project (Grant Agreement no. 951911).

References

- Atanasova, P.; Simonsen, J. G.; Lioma, C.; and Augenstein, I. 2020. A diagnostic study of explainability techniques for text classification. *arXiv preprint arXiv:2009.13295*.
- Belinkov, Y.; and Glass, J. 2019. Analysis Methods in Neural Language Processing: A Survey. *TACL*, 7: 49–72.
- Camburu, O.-M.; Rocktäschel, T.; Lukasiewicz, T.; and Blunsom, P. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.

- Carton, S.; Rathore, A.; and Tan, C. 2020. Evaluating and characterizing human rationales. *arXiv preprint arXiv:2010.04736*.
- Church, K. W.; and Hestness, J. 2019. A survey of 25 years of evaluation. *Natural Language Engineering*, 25(6): 753–767.
- Cliniciu, M.; Eshghi, A.; and Hastie, H. 2021. A study of automatic metrics for the evaluation of natural language explanations. *arXiv preprint arXiv:2103.08545*.
- Costa, F.; Ouyang, S.; Dolog, P.; and Lawlor, A. 2018. Automatic generation of natural language explanations. In *IUI*, 1–2.
- Danilevsky, M.; Qian, K.; Aharonov, R.; Katsis, Y.; Kawas, B.; and Sen, P. 2020. A survey of the state of explainable AI for natural language processing. *arXiv preprint arXiv:2010.00711*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*, 4171–4186.
- DeYoung, J.; Jain, S.; Rajani, N. F.; Lehman, E.; Xiong, C.; Socher, R.; and Wallace, B. C. 2019. ERASER: A benchmark to evaluate rationalized NLP models. *arXiv preprint arXiv:1911.03429*.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Dua, D.; Gottumukkala, A.; Talmor, A.; Gardner, M.; and Singh, S. 2019. Comprehensive Multi-Dataset Evaluation of Reading Comprehension. In *2nd Workshop on Machine Reading for Question Answering*, 147–153. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Dua, D.; Singh, S.; and Gardner, M. 2020. Benefits of Intermediate Annotations in Reading Comprehension. In *ACL*, 5627–5634. ACL.
- Eisenstein, J.; O’Connor, B.; Smith, N. A.; and Xing, E. P. 2014. Diffusion of lexical change in social media. *PLoS one*, 9(11): e113114.
- Garg, S.; Perot, V.; Limtiaco, N.; Taly, A.; Chi, E. H.; and Beutel, A. 2019. Counterfactual fairness in text classification through robustness. In *Conference on AI, Ethics, and Society*, 219–226.
- Glass, A.; McGuinness, D. L.; and Wolverton, M. 2008. Toward Establishing Trust in Adaptive Agents. *IUI*.
- Gunning, D.; Stefik, M.; Choi, J.; Miller, T.; Stumpf, S.; and Yang, G.-Z. 2019. XAI—Explainable artificial intelligence. *Science Robotics*, 4(37): eaay7120.
- Gururangan, S.; Swayamdipta, S.; Levy, O.; Schwartz, R.; Bowman, S.; and Smith, N. A. 2018. Annotation Artifacts in Natural Language Inference Data. In *NAACL*, 107–112.
- Guzman, E. M.; Schlegel, V.; and Batista-Navarro, R. 2022. RaFoLa: A Rationale-Annotated Corpus for Detecting Indicators of Forced Labour. *arXiv preprint arXiv:2205.02684*.
- Hamari, J.; Koivisto, J.; and Sarsa, H. 2014. Does gamification work? - A literature review of empirical studies on gamification. *Annual Hawaii International Conference on System Sciences*, 3025–3034.
- Han, X.; Wallace, B. C.; and Tsvetkov, Y. 2020. Explaining Black Box Predictions and Unveiling Data Artifacts through Influence Functions. *arXiv preprint arXiv:2005.06676*.
- Hase, P.; and Bansal, M. 2020. Evaluating explainable AI: Which algorithmic explanations help users predict model behavior? *arXiv preprint arXiv:2005.01831*.
- Hase, P.; Zhang, S.; Xie, H.; and Bansal, M. 2020. Leakage-adjusted simulatability: Can models generate non-trivial explanations of their behavior in natural language? *arXiv preprint arXiv:2010.04119*.
- Hendricks, L. A.; Hu, R.; Darrell, T.; and Akata, Z. 2018. Generating counterfactual explanations with natural language. *arXiv preprint arXiv:1806.09809*.
- Jansen, P.; Wainwright, E.; Marmostein, S.; and Morrison, C. 2018. WorldTree: A Corpus of Explanation Graphs for Elementary Science Questions supporting Multi-hop Inference. In *LREC*.
- Kiela, D.; Bartolo, M.; Nie, Y.; Kaushik, D.; Geiger, A.; Wu, Z.; Vidgen, B.; Prasad, G.; Singh, A.; Ringshia, P.; Ma, Z.; Thrush, T.; Riedel, S.; Waseem, Z.; Stenetorp, P.; Jia, R.; Bansal, M.; Potts, C.; and Williams, A. 2021. Dynabench: Rethinking Benchmarking in NLP. In *NAACL*, 4110–4124.
- Kocielnik, R.; Amershi, S.; and Bennett, P. N. 2019. Will you accept an imperfect ai? exploring designs for adjusting end-user expectations of ai systems. In *CHI*, 1–14.
- Kowsari, K.; Jafari Meimandi, K.; Heidarysafa, M.; Mendu, S.; Barnes, L.; and Brown, D. 2019. Text classification algorithms: A survey. *Information*, 10(4): 150.
- Krishna, K.; Roy, A.; and Iyyer, M. 2021. Hurdles to Progress in Long-form Question Answering. *NAACL*, 4940–4957.
- Lage, I.; Chen, E.; He, J.; Narayanan, M.; Kim, B.; Gershman, S. J.; and Doshi-Velez, F. 2019. Human evaluation of models built for interpretability. In *Conference on Human Computation and Crowdsourcing*, volume 7, 59–67.
- Lei, T.; Barzilay, R.; and Jaakkola, T. 2016. Rationalizing neural predictions. *arXiv preprint arXiv:1606.04155*.
- Lertvittayakumjorn, P.; and Toni, F. 2019. Human-grounded evaluations of explanation methods for text classification. *arXiv preprint arXiv:1908.11355*.
- Linzen, T. 2020. How Can We Accelerate Progress Towards Human-like Linguistic Generalization? In *ACL*, 5210–5217.
- Liu, H.; Yin, Q.; and Wang, W. Y. 2018. Towards explainable NLP: A generative explanation framework for text classification. *arXiv preprint arXiv:1811.00196*.
- Ma, Q.; Wei, J. T. Z.; Bojar, O.; and Graham, Y. 2019. Results of the WMT19 Metrics Shared Task: Segment-Level and Strong MT Systems Pose Big Challenges. In *WMT*, 62–90.
- McCoy, T.; Pavlick, E.; and Linzen, T. 2019. Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. In *ACL*, 3428–3448.
- Minaee, S.; Kalchbrenner, N.; Cambria, E.; Nikzad, N.; Chenaghlu, M.; and Gao, J. 2021. Deep learning-based text classification: a comprehensive review. *ACM Computing Surveys (CSUR)*, 54(3): 1–40.

- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, D.; and Gebru, T. 2019. Model Cards for Model Reporting. In *FAT*.
- Mohseni, S.; Block, J. E.; and Ragan, E. 2021. Quantitative evaluation of machine learning explanations: a human-grounded benchmark. In *IUI*, 22–31.
- Mohseni, S.; Block, J. E.; and Ragan, E. D. 2018. A human-grounded evaluation benchmark for local explanations of machine learning. *arXiv preprint arXiv:1801.05075*.
- Nguyen, D. 2018. Comparing automatic and human evaluation of local explanations for text classification. In *NAACL*, 1069–1078.
- Nie, Y.; Williams, A.; Dinan, E.; Bansal, M.; Weston, J.; and Kiela, D. 2020. Adversarial NLI: A New Benchmark for Natural Language Understanding. In *ACL*, 4885–4901.
- Otter, D. W.; Medina, J. R.; and Kalita, J. K. 2020. A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2): 604–624.
- Paleja, R.; Ghuy, M.; Arachchige, N. R.; Jensen, R.; and Gombolay, M. 2021. The Utility of Explainable AI in Ad Hoc Human-Machine Teaming. In *NeurIPS*.
- Raamkumar, A. S.; Tan, S. G.; Wee, H. L.; et al. 2020. Use of health belief model-based deep learning classifiers for covid-19 social media content to examine public perceptions of physical distancing: Model development and case study. *JMIR public health and surveillance*, 6(3): e20493.
- Rajagopal, D.; Tandon, N.; Clark, P.; Dalvi, B.; and Hovy, E. 2020. What-if I ask you to explain: Explaining the effects of perturbations in procedural text. In *EMNLP Findings*, 3345–3355.
- Rajani, N. F.; McCann, B.; Xiong, C.; and Socher, R. 2019. Explain yourself! leveraging language models for common-sense reasoning. *arXiv preprint arXiv:1906.02361*.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *EMNLP*, 2383–2392.
- Ribeiro, M. T.; Wu, T.; Guestrin, C.; and Singh, S. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *ACL*, 4902–4912.
- Ross, A. S.; Hughes, M. C.; and Doshi-Velez, F. 2017. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*.
- Samek, W.; Binder, A.; Montavon, G.; Lapuschkin, S.; and Müller, K.-R. 2016. Evaluating the visualization of what a deep neural network has learned. *IEEE Transactions on Neural Networks and Learning Systems*, 28(11): 2660–2673.
- Sawhney, R.; Manchanda, P.; Mathur, P.; Shah, R.; and Singh, R. 2018. Exploring and learning suicidal ideation connotations on social media with deep learning. In *9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 167–175.
- Schlegel, V.; Nenadic, G.; and Batista-Navarro, R. 2020. Beyond Leaderboards: A survey of methods for revealing weaknesses in Natural Language Inference data and models. *arXiv preprint arXiv:2005.14709*.
- Sellam, T.; Das, D.; and Parikh, A. P. 2020. BLEURT: Learning Robust Metrics for Text Generation. In *ACL*, 7881–7892.
- Tadesse, M. M.; Lin, H.; Xu, B.; and Yang, L. 2019. Detection of suicide ideation in social media forums using deep learning. *Algorithms*, 13(1): 7.
- Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; and Mittal, A. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *NAACL*, 809–819.
- Wallace, E.; Feng, S.; Kandpal, N.; Gardner, M.; and Singh, S. 2019a. Universal Adversarial Triggers for Attacking and Analyzing NLP. In *EMNLP*, 2153–2162.
- Wallace, E.; Rodriguez, P.; Feng, S.; Yamada, I.; and Boyd-Graber, J. 2019b. Trick Me If You Can: Human-in-the-Loop Generation of Adversarial Examples for Question Answering. *TACL*, 7: 387–401.
- Wang, A.; Pruksachatkun, Y.; Nangia, N.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. *arXiv preprint arXiv:1905.00537*.
- Wang, A.; Singh, A. A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Weston, J.; Bordes, A.; Chopra, S.; Rush, A. M.; van Merriënboer, B.; Joulin, A.; and Mikolov, T. 2015. Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks. *arXiv preprint arXiv:1502.05698*.
- Wiseman, S.; Shieber, S. M.; and Rush, A. M. 2017. Challenges in Data-to-Document Generation. In *EMNLP*, 2253–2263.
- Zaidan, O.; Eisner, J.; and Piatko, C. 2007. Using “Annotator Rationales” to Improve Machine Learning for Text Categorization - ACL Anthology. In *NAACL*, 260–267. Association for Computational Linguistics.
- Zellers, R.; Holtzman, A.; Clark, E.; Qin, L.; Farhadi, A.; and Choi, Y. 2020. Evaluating Machines by their Real-World Language Use. *arXiv preprint arXiv:2004.03607*.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019. BERTScore: Evaluating Text Generation with BERT. *arXiv preprint, arXiv:1904.09675*.