

An Investigation of Language Model Interpretability via Sentence Editing

Samuel Stevens

The Ohio State University
stevens.994@osu.edu

Yu Su

The Ohio State University
su.809@osu.edu

Abstract

Pre-trained language models (PLMs) like BERT are being used for almost all language-related tasks, but interpreting their behavior still remains a significant challenge and many important questions remain largely unanswered. In this work, we re-purpose a sentence editing dataset, where faithful high-quality human rationales can be automatically extracted and compared with extracted model rationales, as a new testbed for interpretability. This enables us to conduct a systematic investigation on an array of questions regarding PLMs’ interpretability, including the role of pre-training procedure, comparison of rationale extraction methods, and different layers in the PLM. The investigation generates new insights, for example, contrary to the common understanding, we find that attention weights correlate well with human rationales and work better than gradient-based saliency in extracting model rationales. Both the dataset and code are available at <https://github.com/samuelstevens/sentence-editing-interpretability> to facilitate future interpretability research.

1 Introduction

Pre-trained language models (PLMs) (Devlin et al., 2019; Liu et al., 2019; Beltagy et al., 2019) are pervasively used in language-related tasks, but interpreting their predictions is notoriously difficult because of their parameters’ complex interdependencies. Given a specific prediction, we want to know *why* a model made that decision, both to further improve performance and to use the model in high-stakes scenarios such as healthcare or bank loan approvals where interpretability is important. This has motivated efforts in extracting model explanations, typically in the form of *rationales*, i.e., subsets of the original input that support a decision (Zaidan et al., 2007). Attention heatmaps (Xu et al., 2015) and gradient-based saliency maps (Simonyan et al., 2014) are common extraction methods.

The algorithm ~~described~~described in the previous sections has several advantages.

The algorithm ~~described~~→**described** in the previous sections has several advantages.

However, we ~~must note that~~ we ~~still~~ have no means of deciding which documents out of `_MATH_` deserve to be in `_MATH_` and `_MATH_`, respectively.

However, we ~~must note that~~ we ~~still~~ have no means of deciding which documents out of `_MATH_` deserve to be in `_MATH_` and `_MATH_`, respectively.

Figure 1: Two “need edit” examples from AESW in the original data format and a human-readable format. The first example (a) has a spelling error “described” and the second (b) is edited for concision.

There have been efforts on developing datasets for interpretability research, for example, the recent ERASER benchmark (DeYoung et al., 2020). However, the majority of ERASER tasks use human rationales highlighted by a different annotator after the original labeling process. Such rationales are not necessarily *faithful*; a rationale highlighted by the second annotator may not have been actually used by the first annotator while labeling. Manual rationale labeling is also difficult and time-consuming; of the six datasets in the ERASER benchmark, only one has more than 200 examples.

Our first contribution is the realization that AESW (Automatic Evaluation of Scientific Writing; Daudaravicius et al., 2016), a sentence editing dataset, contains *thousands* of *faithful human rationales* that can be automatically re-purposed for interpretability research. See Figure 1 for examples. This provides a new, large-scale dataset with truly faithful human rationales for interpretability questions surrounding model rationales.

Our second contribution is investigating multiple factors in PLM rationale plausibility. More plausible rationales are valuable in human-in-the-

loop systems where humans use model rationales to make a final decision. We compare (1) pre-training procedures, (2) attention weight- and input gradient-based methods of extracting model rationales, (3) correlation between model rationale plausibility and model confidence, and (4) differences in transformer layers. While previous work (Jain and Wallace, 2019; Serrano and Smith, 2019) has shown that attention weights are not always faithful, we find that they correlate with human rationales better than gradient-based methods.

2 Related Work

Human rationales (as defined by Zaidan et al., 2007) are subsets of input highlighted by human annotators as evidence to support a decision. The same annotator labeling an example might also highlight their rationale (Khashabi et al., 2018; Thorne et al., 2018). In other cases, rationales are collected for an existing dataset by different annotators (Zaidan et al., 2008; Camburu et al., 2018; Rajani et al., 2019). As previously stated, such rationales may not be faithful. Rationale length can vary from sub-sentence spans (Talmor et al., 2019) to multiple sentences (Lehman et al., 2019).

Model rationales can be produced as an explicit training objective (Zaidan et al., 2008) or extracted as a post-hoc explanation. Post-hoc methods typically assign token-level importance scores: attention weights are often used in attention-based models (Bahdanau et al., 2016), gradient-based explanations are typical for differentiable models (Denil et al., 2015; Shrikumar et al., 2017), and LIME is a model-agnostic method (Ribeiro et al., 2016). We follow work using BERT’s attention (Clark et al., 2019; Kovaleva et al., 2019) to extract rationales.

A model rationale is evaluated on faithfulness (if it is actually used to make a decision) and plausibility (if it is easily understood by humans). Faithfulness can be measured by perturbing inputs marked as evidence and measuring change in outputs (Jain and Wallace, 2019; Serrano and Smith, 2019). Plausibility can be measured through user studies, wherein users are given a model rationale and asked either to predict the model’s decision (Kim et al., 2016) or to rate rationale understandability (Nguyen, 2018; Ehsan et al., 2018, 2019; Strout et al., 2019). Rationale plausibility can also be measured by similarity to human rationales (DeYoung et al., 2020), but this requires faithful human rationales. We use similarity to evaluate

rationale plausibility because we gather faithful human rationales from sentence editing annotations.

3 Proposed Task

We propose re-purposing the AESW classification task for measuring model interpretability. We gather examples from AESW from which we can automatically extract faithful and sufficient human rationales, and then use said rationales to investigate factors in PLM rationale plausibility, specifically BERT (Devlin et al., 2019) and its variants. It is worth noting that our rationale dataset can be used for other interpretability topics such as training with rationales or evaluating rationale faithfulness; we will focus on plausibility considering the scope of this paper.

3.1 Human Rationales

Human rationales are substrings used as evidence for a decision (Zaidan et al., 2007). Faithful and sufficient (enough evidence to justify a decision) human rationales can be used as gold labels for evaluating model rationale plausibility.

The original AESW task is to predict if a sentence from a scientific paper needs editing. Daudaravicius et al. extract spans of a sentence before and after professional editing¹ as deleted (``) or inserted (`<ins>`) and provide 1.1M training, 140K validation and 140K testing examples. Sentences without changes are assumed to not require editing.

We exploit the data format to automatically extract faithful and sufficient human rationales. Delete text (text between `` tags) is always a *faithful* rationale (such text is a source of the error). For edits with arbitrary `` text alone and `<ins>` text, the `` alone is not always a *sufficient* rationale to justify “need edit”. Consider a sentence where a verb is incorrectly conjugated and replaced with the correctly conjugated verb. The incorrectly conjugated verb is not sufficient to decide that the sentence needs editing; the surrounding context is required. To find edits where `` text *is* always a sufficient rationale, we use two criteria:

1. A spelling error is corrected. (*spelling error*)
2. Text is only deleted, not added. (*deleted text*)

Spelling errors are always a sufficient rationale to justify editing a sentence (see Figure 1a). In edits with no insertions, removing the `` text leads to an error-free sentence, so the `` text

¹Native English speaking editors working at VTeX.

is sufficient explanation for editing (see Figure 1b). These criteria lead to both a simple, lexical task (*spelling error*) and a more challenging semantic task (*deleted text*) and constitute a wider range of challenges for future interpretability research. We extract faithful and sufficient human rationales for 1,321 *spelling error* edits and 6,741 *deleted text* edits from the validation set of AESW.²

3.2 Model Rationales

Model rationales are substrings provided by a model as evidence for a decision. Given a model, an example x_i and a prediction y_i , we extract three model rationales.

First, we use attention maps (Xu et al., 2015; Kovaleva et al., 2019) to rank word relevance. We measure the total attention weight from BERT’s final layer’s initial [CLS] token to each token t_j across H attention heads. Then we add those totals together for each token t_j in a word w :

$$\text{score}(w) = \sum_{t_j \in w} \sum_{h=1}^H \text{Attn}_{h_j}([\text{CLS}] \rightarrow t_j). \quad (1)$$

Second, we use gradient-based saliency (specifically gradient×input; Denil et al., 2015; Shrikumar et al., 2017) to rank word relevance. We calculate a saliency score for each token t_j in x_i . First, change in model output with respect to t_j ’s input embedding $\nabla_{e(t_j)} f_{y_i}(x_i)$ captures the sensitivity to token t . The dot product with $e(t_j)$ is then a scalar measure of each token’s marginal impact on the model prediction (Han et al., 2020). Finally, we again compute a word-level score by summing over each token t_j in word w :

$$\text{score}(w) = \sum_{t_j \in w} \nabla_{e(t_j)} f_{y_i}(x_i) \cdot e(t_j) \quad (2)$$

In contrast to attention, gradient×input faithfully measures the marginal effect of each input token on the prediction (Bastings and Filippova, 2020). Finally, we extract a third set of rankings using gradient×input’s magnitude to rank words.

3.3 Evaluation

We evaluate extracted rationale plausibility using similarity to human rationales (DeYoung et al.,

²More details on extracting human rationales, as well as additional examples, can be found in Appendix C

2020). We use the continuous word scores generated in the previous section to rank relevance, then use mean reciprocal rank as an evaluation metric.³

4 Experiments

To demonstrate the utility of the AESW task for interpretability research, we present four experiments, each with the goal of understanding factors in PLM rationale plausibility.

For our experiments, we add a linear layer and sigmoid activation function on top of the [CLS] token representation, fine-tune BERT-base end-to-end on the AESW training set using the original training objective (classify a sentence as “need edit” or “no edit”) and use validation loss to tune hyperparameters. We do not add any interpretability or rationale-related objectives. The classification F1 for the models on *spelling error* and *deleted text* edits is 82.0 and 74.0, respectively.⁴ We extract and evaluate model rationales on all *spelling error* and *deleted text* edits, even edits that a model does not correctly predict as “need edit.”

Does pre-training procedure affect rationale plausibility? We are interested in how pre-training affects plausibility after fine-tuning (to the best of our knowledge, this is a previously unexplored topic). We compare BERT and two variants, RoBERTa (Liu et al., 2019) and SciBERT (Beltagy et al., 2019). These models have the same architecture but differ in pre-training: RoBERTa is pre-trained for 5x longer with 10x more data than BERT, and SciBERT is pre-trained on a corpus of academic papers. We extract model rationales using attention weights and evaluate their plausibility.

As seen in Figure 2, RoBERTa and BERT generate nearly equally plausible rationales despite differences in pre-training corpus size. We hypothesize that SciBERT generates less plausible rationales because it encodes “need edit” representations in earlier layers (rather than in the final layer), then attends to [SEP] as a no-op in later layers, as proposed in Clark et al. (2019) and Kobayashi et al. (2020) and further confirmed in the subsequent experiments.

³We performed our analyses with additional, similar metrics; our main observations remain consistent independent of the choice of metric. Appendix D contains additional details and complete results.

⁴Appendix B contains more details on fine-tuning procedure and results.

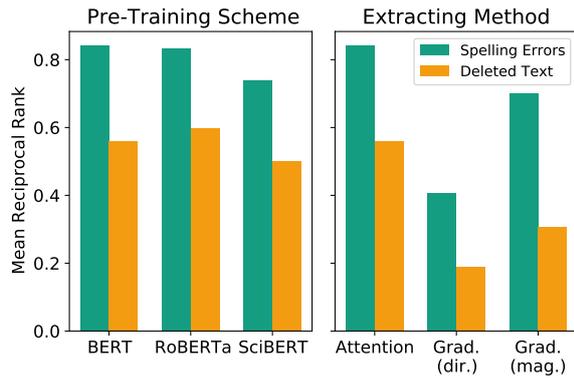


Figure 2: **Left:** BERT, RoBERTa and SciBERT’s attention weight rationale plausibility. **Right:** BERT’s attention weight, gradient \times input and $|\text{gradient} \times \text{input}|$ rationale plausibility (see Section 3.2).

Do attention weights or input gradients produce better rationales? In contrast to attention weights, gradient \times input scores are naturally faithful with respect to individual feature importance (Jain and Wallace, 2019; Bastings and Filippova, 2020). However, attention weights can represent word relevance *in context*, potentially leading to more plausible rationales. We extract and evaluate rationales using attention weights and the two gradient \times input methods described in Section 3.2.

Figure 2 shows that attention-based rationales are more plausible and that the difference is more pronounced on *deleted text* edits. Using $|\text{gradient} \times \text{input}|$ (right-most) also shows improvements over directional gradient \times input (middle-right), in contrast to Han et al. (2020).

Are plausibility and confidence correlated? We are curious if BERT is more confident in its classification decisions when it attends to the original evidence used by editors. We quantify BERT’s classification confidence by treating the sigmoid activation function’s value as the probability of a “need edit” decision. We only consider examples that BERT correctly classifies as “need edit” and calculate BERT’s mean confidence. We find that when human rationales are BERT’s most attended-to words, it is 11.7% and 11.9% more confident in its predictions for *spelling error* and *deleted text* edits, respectively.

One possible explanation for such a correlation is that easy-to-classify examples also have easy-to-identify rationales. Because agreement with human rationales is not the training objective, however, we believe this correlation suggests that BERT learns to classify edits similarly to humans.

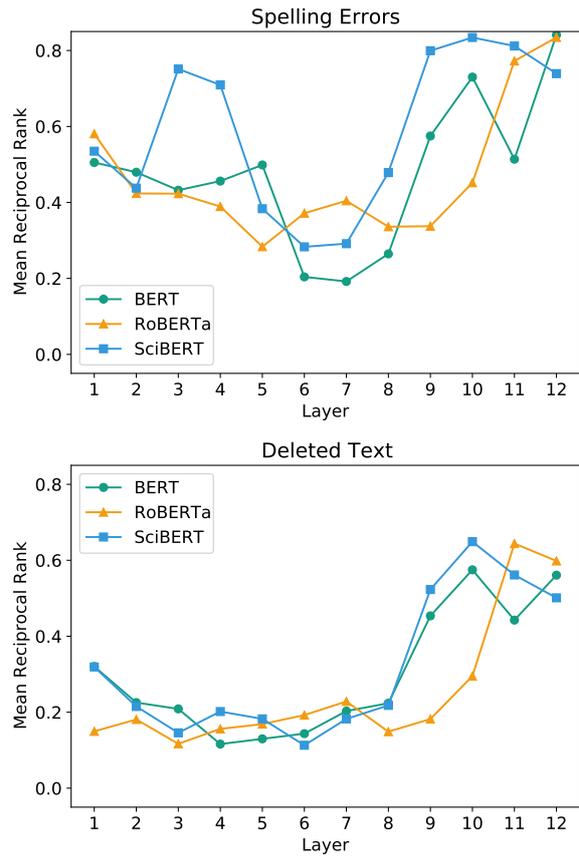


Figure 3: Mean reciprocal rank for each layer (using the *mean* strategy) for each model for *spelling error* and *deleted text* edits.⁴

How does transformer layer affect plausibility? It is widely agreed that BERT’s early layers encode more lexical and phrasal information than other layers (Jawahar et al., 2019; Rogers et al., 2020). We hypothesize that rationales extracted from early layers will be more plausible for *spelling error* than *deleted text* edits because it is a more lexical task. We extract rationales from each layer’s attention weights and measure their plausibility in Figure 3.

The results confirm the hypothesis and show that early layers in BERT models are indeed more lexically-oriented. We also find that SciBERT strongly attends to spelling errors in earlier layers. We believe it is because that, pre-trained on well-formed academic text, SciBERT is not as well exposed to spelling errors as BERT and RoBERTa so it learns to attend to spelling errors in earlier layers during fine-tuning.

⁴BERT shows a decline in plausibility at layer 11 in both edit types because it attends heavily to periods.

5 Conclusion

We re-purpose the AESW task to gather thousands of inherently faithful human rationales and investigate an array of questions regarding PLM interpretability. We find, among other new insights, that attention weights correlate well with human rationales and produce more plausible rationales than input gradients, which is different from existing understanding. Furthermore, we find that BERT is more confident in its predictions when it attends to the same words that a human did, supporting the idea that while attention is not inherently faithful, attention-based models might rely on the same information as humans when making a prediction.

Future work might expand the subset of examples for which human rationales can be automatically extracted, include human rationales during training or evaluate faithful-by-design model rationales on this dataset.

6 Acknowledgements

The authors would like to extend their thanks to the anonymous reviewers for their insightful feedback.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2016. [Neural machine translation by jointly learning to align and translate](#). ArXiv preprint arXiv:1409.0473.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In [Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP](#), pages 149–155, Online. Association for Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-snli: Natural language inference with natural language explanations](#). In [Advances in Neural Information Processing Systems](#), volume 31, pages 9539–9549. Curran Associates, Inc.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In [Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP](#), pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Vidas Daudaravicius, Rafael E. Banchs, Elena Volodina, and Courtney Napoles. 2016. [A report on the automatic evaluation of scientific writing shared task](#). In [Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications](#), pages 53–62, San Diego, CA. Association for Computational Linguistics.
- Misha Denil, Alban Demiraj, and Nando de Freitas. 2015. [Extraction of salient sentences from labelled documents](#). ArXiv preprint arXiv:1412.6815.
- 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.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. [ERASER: A benchmark to evaluate rationalized NLP models](#). In [Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics](#), pages 4443–4458, Online. Association for Computational Linguistics.
- Upol Ehsan, Brent Harrison, Larry Chan, and Mark O. Riedl. 2018. [Rationalization: A neural machine translation approach to generating natural language explanations](#). In [Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES ’18](#), page 81–87, New York, NY, USA. Association for Computing Machinery.
- Upol Ehsan, Pradyumna Tambwekar, Larry Chan, Brent Harrison, and Mark O. Riedl. 2019. [Automated rationale generation: A technique for explainable ai and its effects on human perceptions](#). In [Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI ’19](#), page 263–274, New York, NY, USA. Association for Computing Machinery.
- 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.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). 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 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In [Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics](#), pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. [Looking beyond the surface: A challenge set for reading comprehension over multiple sentences.](#) In [Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 \(Long Papers\)](#), pages 252–262, New Orleans, Louisiana. Association for Computational Linguistics.
- Been Kim, Rajiv Khanna, and Oluwasanmi Koyejo. 2016. [Examples are not enough, learn to criticize! criticism for interpretability.](#) In [Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16](#), page 2288–2296, Red Hook, NY, USA. Curran Associates Inc.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms.](#) In [Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing \(EMNLP\)](#), pages 7057–7075, Online. Association for Computational Linguistics.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT.](#) 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 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Eric Lehman, Jay DeYoung, Regina Barzilay, and Byron C. Wallace. 2019. [Inferring which medical treatments work from reports of clinical trials.](#) 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 3705–3717, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Dong Nguyen. 2018. [Comparing automatic and human evaluation of local explanations for text classification.](#) In [Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 \(Long Papers\)](#), pages 1069–1078, New Orleans, Louisiana. Association for Computational Linguistics.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain yourself! leveraging language models for commonsense reasoning.](#) In [Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics](#), pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Marco Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. [“why should I trust you?”: Explaining the predictions of any classifier.](#) In [Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations](#), pages 97–101, San Diego, California. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in bertology: What we know about how bert works.](#) [Transactions of the Association for Computational Linguistics](#), 8:842–866.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In [Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics](#), pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. [Learning important features through propagating activation differences.](#) In [Proceedings of the 34th International Conference on Machine Learning](#), volume 70 of [Proceedings of Machine Learning Research](#), pages 3145–3153, International Convention Centre, Sydney, Australia. PMLR.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. [Deep inside convolutional networks: Visualising image classification models and saliency maps.](#) ArXiv preprint arXiv:1312.6034.
- Julia Strout, Ye Zhang, and Raymond Mooney. 2019. [Do human rationales improve machine explanations?](#) In [Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP](#), pages 56–62, Florence, Italy. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge.](#) 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 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018.

[FEVER: a large-scale dataset for fact extraction and VERification](#). In [Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 \(Long Papers\)](#), pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#). ArXiv preprint arXiv:1910.03771.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In [Proceedings of the 32nd International Conference on Machine Learning](#), volume 37 of [Proceedings of Machine Learning Research](#), pages 2048–2057, Lille, France. PMLR.

Omar Zaidan, Jason Eisner, and Christine Piatko. 2007. [Using “annotator rationales” to improve machine learning for text categorization](#). In [Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference](#), pages 260–267, Rochester, New York. Association for Computational Linguistics.

Omar F. Zaidan, Jason Eisner, and Christine Piatko. 2008. [Machine learning with annotator rationales to reduce annotation cost](#). In [Proceedings of the NIPS*2008 Workshop on Cost Sensitive Learning](#).

A AESW Details

The original AESW task is to classify a sentence from an academic paper as “need edit” or “no edit”. There is no special markup attached to the sentence when it is given as input to a model. Task participants receive the training and validation sets with `` and `<ins>` tokens for model development, and Daudaravicius et al. provided an automatic test set evaluation through an online portal. Daudaravicius et al. collected these sentences from 9,919 journal articles published by Springer Publishing Company and edited at VTeX. Sentences before and after editing were automatically aligned using a modified diff algorithm.⁵ After the challenge, Daudaravicius et al. released the dataset in its entirety (all three datasets with all tokens/spans included).

B Training Details

The AESW task uses scientific papers written in LaTeX, which contains markup characters that impact sentence meaning. The original authors (Daudaravicius et al.) replace these characters with special tokens, as seen in Table 1. We add these four special tokens (`_MATH_`, `_MATHDISP_`, `_CITE_` and `_REF_`) to the model vocabulary, fine-tuning the word representations during training.

LaTeX Example	Special Token
<code>\beta_{2}</code>	<code>_MATH_</code>
<code>\$\$2 + 3\$\$</code>	<code>_MATHDISP_</code>
<code>\cite{google2018}</code>	<code>_CITE_</code>
<code>\ref{tab:results}</code>	<code>_REF_</code>

Table 1: Special tokens found in the original AESW data that should not be split further into bytes/tokens.

We train all models for a maximum of 30 epochs with a patience of 5 on a single Tesla P100 GPU.

All models (BERT⁶, SciBERT⁷, RoBERTA⁸) are based on their HuggingFace (Wolf et al., 2020) implementations.

We list all the key hyperparameters and tuning bounds for reproducibility in Table 3. Our final results for all three BERT-based models along with

⁵More details can be found in (Daudaravicius et al., 2016)

⁶https://huggingface.co/transformers/v3.0.2/model_doc/bert.html#bertforsequenceclassification

⁷<https://github.com/allenai/scibert#pytorch-huggingface-models>

⁸https://huggingface.co/transformers/v3.0.2/model_doc/roberta.html#robertaforsequenceclassification

the top three models from (Daudaravicius et al., 2016) can be found in Table 2. Additionally, we will release code and instructions for reproducing our results.

Model	Dev F1	Test F1
CNN+LSTM	–	0.628
CNN	–	0.611
SVM	–	0.555
BERT _{base}	0.654	0.666
RoBERTA _{base}	0.661	0.670
SciBERT _{scivocab}	0.658	0.668

Table 2: Performance on the original AESW sentence classification task. Dev set results are not available for models reported in Daudaravicius et al. (2016).

C Gathering Human Rationales

To find *spelling error* edits, we look for sentences with a deleted, misspelled word followed by an inserted, correctly spelled word. The first two examples in Figure 4 are examples of *spelling error* edits, while the third is not.

To find *deleted text* edits, we look for sentences where text is removed but not added. The fourth and fifth examples in Figure 4 are examples of *deleted text* edits.

D Evaluating Model Rationales

Although the main text presents our analyses using mean reciprocal rank, we performed our analyses with multiple metrics. Here we provide specific definitions for our metrics and our complete results for all models with every metric.

Mean reciprocal rank Because human rationales can be made up of multiple words, we need to modify mean reciprocal ranking. For a single sentence, a model’s ordered ranking of words M and an unordered human rationale H :

1. Find the top ranked word w in M from H and record the rank.
2. Remove w from M and H .
3. Repeat until H is empty.
4. Use the reciprocal of the mean rank.

If we did not remove words from M , a perfect score for a rationale with multiple words would be impossible. Consider the sentence “The boy ate **the found his** ball.” If the rationale ranked ‘ate’ and ‘the’ as most important, the mean rationale rank would be $(1 + 2)/2 = 1.5$ and the rationale’s reciprocal rank would be $1/1.5 = 0.66$, despite a perfect rationale.

Model	Hyperparameters	Hyperparameter bounds
BERT_{base}	learning rate: 1×10^{-6}	learning rate: (2×10^{-7} , 1×10^{-6} , 2×10^{-5} , 1×10^{-4})
	batch size: 32 model: bert-base-uncased vocab size: 30526 (normally 30522)	
RoBERTa_{base}	learning rate: 1×10^{-6}	learning rate: (1×10^{-6})
	batch size: 32 model: roberta-base vocab size: 50269 (normally 50265)	
SciBERT	learning rate: 1×10^{-6}	learning rate: (1×10^{-6})
	batch size: 32 model: allenai/scibert_scivocab_uncased vocab size: 31094 (normally 31090)	

Table 3: Hyperparameter options for each model. Note that each model had 4 special tokens added to the vocabulary. BERT was fine-tuned first. Because of compute limitations, RoBERTa and SciBERT were both fine-tuned using the same hyperparameters as the optimal BERT configuration (learning rate of (1×10^{-6})).

We take the mean reciprocal ranking across all examples to evaluate rationales.

Mean area under precision-recall curve Using the model relevance scores for words in a sentence, we adjust the threshold for classifying a word as part of the rationale, calculate a precision-recall curve and measure the area underneath. We take the mean AUPRC across all examples.

Mean top 1 match We score a rationale as 1 if the rationale’s top ranked word is the human rationale and 0 otherwise. This means that all multiple-word rationales are automatically not a match and scored as 0. We take the mean of these scores across all examples.

Our main observations are consistent across metrics: Table 4 contains all results for models on *spelling error* edits and Table 5 contains all results for models on *deleted text* edits.

Model	F1	Method	Classification	Mean Recip.	AUPRC	Top 1 Match
BERT	82.1	Attention	All	0.841	0.115	0.756
BERT	100.0	Attention	Correct	0.895	0.081	0.824
BERT	0.0	Attention	Wrong	0.717	0.194	0.602
BERT	82.1	gradient×input	All	0.407	0.332	0.325
BERT	100.0	gradient×input	Correct	0.401	0.335	0.322
BERT	0.0	gradient×input	Wrong	0.422	0.325	0.331
BERT	82.1	gradient × input	All	0.701	0.201	0.579
BERT	100.0	gradient × input	Correct	0.730	0.184	0.615
BERT	0.0	gradient × input	Wrong	0.635	0.241	0.498
RoBERTa	81.4	Attention	All	0.834	0.118	0.739
RoBERTa	100.0	Attention	Correct	0.886	0.085	0.811
RoBERTa	0.0	Attention	Wrong	0.721	0.192	0.581
RoBERTa	81.4	gradient×input	All	0.310	0.380	0.231
RoBERTa	100.0	gradient×input	Correct	0.303	0.383	0.229
RoBERTa	0.0	gradient×input	Wrong	0.324	0.375	0.236
RoBERTa	81.4	gradient × input	All	0.675	0.219	0.546
RoBERTa	100.0	gradient × input	Correct	0.707	0.201	0.583
RoBERTa	0.0	gradient × input	Wrong	0.605	0.259	0.463
SciBERT	86.8	Attention	All	0.739	0.203	0.577
SciBERT	100.0	Attention	Correct	0.830	0.138	0.707
SciBERT	0.0	Attention	Wrong	0.442	0.415	0.152
SciBERT	86.8	gradient×input	All	0.475	0.296	0.399
SciBERT	100.0	gradient×input	Correct	0.467	0.296	0.401
SciBERT	0.0	gradient×input	Wrong	0.501	0.295	0.392
SciBERT	86.8	gradient × input	All	0.670	0.214	0.557
SciBERT	100.0	gradient × input	Correct	0.698	0.197	0.591
SciBERT	0.0	gradient × input	Wrong	0.575	0.268	0.447

Table 4: Results for all three models for *spelling error* edits across all methods of selecting rationales (attention, gradient×input and |gradient × input|) and all metrics to evaluate rationales (mean reciprocal ranking, AUPRC and mean top 1 match).

Model	F1	Method	Classification	Mean Recip.	AUPRC	Top 1 Match
BERT	74.0	Attention	All	0.561	0.272	0.411
BERT	100.0	Attention	Correct	0.680	0.211	0.528
BERT	0.0	Attention	Wrong	0.391	0.359	0.244
BERT	74.0	gradient×input	All	0.189	0.447	0.074
BERT	100.0	gradient×input	Correct	0.195	0.442	0.089
BERT	0.0	gradient×input	Wrong	0.181	0.454	0.054
BERT	74.0	gradient × input	All	0.306	0.416	0.130
BERT	100.0	gradient × input	Correct	0.354	0.399	0.166
BERT	0.0	gradient × input	Wrong	0.238	0.440	0.078
RoBERTa	74.8	Attention	All	0.598	0.263	0.426
RoBERTa	100.0	Attention	Correct	0.662	0.230	0.487
RoBERTa	0.0	Attention	Wrong	0.503	0.312	0.335
RoBERTa	74.8	gradient×input	All	0.168	0.456	0.058
RoBERTa	100.0	gradient×input	Correct	0.156	0.458	0.054
RoBERTa	0.0	gradient×input	Wrong	0.187	0.452	0.063
RoBERTa	74.8	gradient × input	All	0.321	0.404	0.154
RoBERTa	100.0	gradient × input	Correct	0.367	0.384	0.194
RoBERTa	0.0	gradient × input	Wrong	0.252	0.434	0.094
SciBERT	73.9	Attention	All	0.501	0.327	0.304
SciBERT	100.0	Attention	Correct	0.605	0.270	0.414
SciBERT	0.0	Attention	Wrong	0.353	0.406	0.148
SciBERT	73.9	gradient×input	All	0.226	0.426	0.111
SciBERT	100.0	gradient×input	Correct	0.260	0.410	0.145
SciBERT	0.0	gradient×input	Wrong	0.178	0.450	0.062
SciBERT	73.9	gradient × input	All	0.322	0.400	0.162
SciBERT	100.0	gradient × input	Correct	0.356	0.387	0.189
SciBERT	0.0	gradient × input	Wrong	0.275	0.418	0.123

Table 5: Results for all three models for *deleted text* edits across all methods of selecting rationales (attention, gradient×input and |gradient × input|) and all metrics to evaluate rationales (mean reciprocal ranking, AUPRC and mean top 1 match).

The algorithm ~~descripted~~described in the previous sections has several advantages.

A spelling error edit; “descripted” is a spelling error and is corrected.

For each energy point, the thereby obtained cross-section values and errors from different experiments have been further averaged according to the ~~weigthed~~weighted average method used by the Particle Data Group _CITE_, including error rescaling by _MATH_ in case of large discrepancy.

A spelling error edit; “weigthed” is a spelling error and is corrected.

And the short notations for the denominators are _MATHDISP_, Furthermore, the following relations are useful to ~~short~~shortcut the expressions: _MATHDISP_.

Not a spelling error edit; “short” is incorrect in this context, but it is not a spelling error.

However, we ~~must note that we~~ still have no means of deciding which documents out of _MATH_ deserve to be in _MATH_ and _MATH_, respectively.

A deleted text edit; text is only removed, while no text is inserted.

Let _MATH_ be the conjugate Holder~~'s~~ function of a Holder~~'s~~ function _MATH_.

A deleted text edit; text is only removed, while no text is inserted.

Figure 4: Additional examples of the two types of edits extracted from the original AESW dataset.