

The Irrationality of Neural Rationale Models

Yiming Zheng Serena Booth Julie Shah Yilun Zhou
MIT CSAIL

yimingz@mit.edu {serenabooth, julie_a_shah, yilun}@csail.mit.edu

Abstract

Neural rationale models are popular for interpretable predictions of NLP tasks. In these, a selector extracts segments of the input text, called *rationales*, and passes these segments to a classifier for prediction. Since the rationale is the only information accessible to the classifier, it is plausibly *defined* as the explanation. Is such a characterization unconditionally correct? In this paper, we argue to the contrary, with both philosophical perspectives and empirical evidence suggesting that rationale models are, perhaps, less rational and interpretable than expected. We call for more rigorous evaluations of these models to ensure desired properties of interpretability are indeed achieved. The code for our experiments is at <https://github.com/yimingz89/Neural-Rationale-Analysis>.

1 Introduction

As machine learning models are increasingly used in high-stakes domains, understanding the reasons for a prediction becomes more important, especially when the model is a black-box such as a neural network. While many *post-hoc* interpretability methods have been developed for models operating on tabular, image, and text data (Simonyan et al., 2013; Ribeiro et al., 2016; Feng et al., 2018), their faithfulness are often questioned (Adebayo et al., 2018; Rudin, 2019; Zhou et al., 2022a).

With no resolution in sight for explaining black box models, *inherently interpretable* models, which self-explain while making decisions, are often favored. Neural rationale models, shown in Figure 1 (top), are the most popular in NLP (Lei et al., 2016; Bastings et al., 2019; Yu et al., 2019; Jain et al., 2020): in them, a selector processes the input text, extracts segments (i.e. *rationale*) from it, and sends *only* the rationale to the predictor. Since the rationale is the only information accessible to the predictor, it arguably serves as the *explanation* for the prediction.

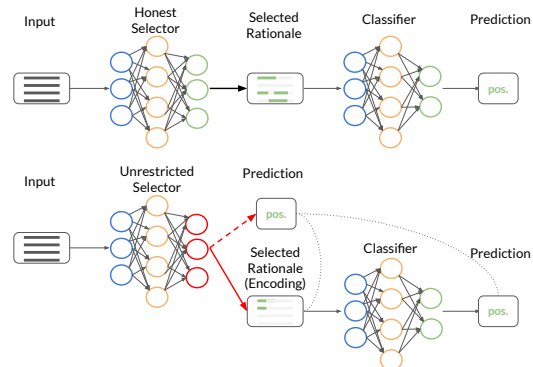


Figure 1: Top: an honest neural rationale model. We seek to understand the selector’s process (the bold arrow), which should select words and phrases as rationale in an unbiased way, leaving the prediction to the classifier which receives this rationale. Bottom: a failure case of neural rationale models. As discussed in Section 3, an unrestricted selector may be able to make its own (relatively accurate) prediction, and “pass” it to the classifier via encoding it in the selected rationale.

While the bottleneck structure defines a causal relationship between rationale and prediction, we caution against equating this structure with inherent interpretability without additional constraints. Notably, if both the selector and the classifier are sufficiently flexible function approximators (e.g. neural networks), the bottleneck structure provides *no* intrinsic interpretability as the selector and classifier may exploit imperceptible messages, as shown in Figure 1 (bottom).

We perform a suite of empirical analyses to demonstrate how rationales lack interpretability. Specifically, we present modes of instability of the rationale selection process under minimal and meaning-preserving sentence perturbations on the Stanford Sentiment Treebank (SST, Socher et al., 2013) dataset. Through a user study, we further show that this instability is poorly understood by people—even those with advanced machine learning knowledge. We find that the exact form of interpretability induced by neural rationale models, if any, is not clear. As a community, we must criti-

cally reflect on the interpretability of these models, and perform rigorous evaluations about any and all claims of interpretability going forward.

2 Related Work

Most interpretability efforts focus on *post-hoc* interpretation. For a specific input, these methods generate an explanation by analyzing model behaviors such as gradient (Simonyan et al., 2013; Sundararajan et al., 2017) or prediction on perturbed (Ribeiro et al., 2016; Lundberg and Lee, 2017) or reduced (Feng et al., 2018) inputs. However, evaluations of these methods highlight various problems. For example, Adebayo et al. (2018) showed that many methods can generate seemingly reasonable explanations even for random neural networks. Zhou et al. (2022a) found that many methods fail to identify features known to be used by the model. Zhou et al. (2022b) share the same principles as us, but also focus on general post-hoc interpretations of arbitrary black-box models, while we focus on neural rationale models.

By contrast, neural rationale models are largely deemed *inherently interpretable* and thus do not require *post-hoc* analysis. At a high level, a model has a selector and a classifier. For an input sentence, the selector first calculates the *rationale* as excerpts of the input, and then the classifier makes a prediction from *only* the rationale. Thus, the rationale is often *defined* as the explanation due to this bottleneck structure. The non-differentiable rationale selection prompts people to train the selector using policy gradient (Lei et al., 2016; Yu et al., 2019) or continuous relaxation (Bastings et al., 2019), or directly use a pre-trained one (Jain et al., 2020).

While rationale models have mostly been subject to less scrutiny, some evaluations have been carried out. Yu et al. (2019) proposed the notions of comprehensiveness and sufficiency for rationales, advocated as standard evaluations in the ERASER (DeYoung et al., 2019) dataset. Zhou et al. (2022a) noted that training difficulty, especially due to policy gradient, leads to selection of words known to not influence the label in the data generative model. Complementing these evaluations and criticisms, we argue from additional angles to be wary of interpretability claims for rationale models, and present experiments showing issues with existing models.

Most related to our work, Jacovi and Goldberg (2020) mention a Trojan explanation and dominant selector as two failure modes of rationale mod-

els. We pinpoint the same root cause of a non-understandable selector in Section 3. However, they favor rationales generated *after* the prediction, while we will argue for rationales being generated *prior* to the prediction. Also, in their discussion of contrastive explanations, their proposed procedure runs the model on out-of-distribution data (sentence with some tokens masked), potentially leading to arbitrary predictions due to extrapolation, a criticism also argued by Hooker et al. (2018).

3 Philosophical Perspectives

In neural rationale models, the classifier prediction causally results from the selector rationale, but does this property automatically equate rationale with explanation? We first present a “failure case.” For a binary sentiment classification, we first train a (non-interpretable) classifier c' that predicts on the whole input. Then we *define* a selector s' that selects the first word of the input if the prediction is positive, or the first two words if the prediction is negative. Finally, we train a classifier c to imitate the prediction of c' but from the rationale. The $c' \rightarrow s' \rightarrow c$ model should achieve best achievable accuracy, since the actual prediction is made by the unrestricted classifier c' with full input access. Can we consider the rationale as explanation? No, because the rationale selection depends on, and is as (non-)interpretable as, the black-box c' . This failure case is shown in Figure 1 (bottom). Recently proposed introspective training (Yu et al., 2019) could not solve this problem either, as the selector can simply output the comprehensive rationale along with the original cue of first one or two words, with only the latter used by the classifier¹. In general, a sufficiently powerful selector can make the prediction at selection time, and then pass this prediction via some encoding in the selected rationale for the classifier to use.

To hide the “bug,” consider now s' selecting the three most positive or negative words in the sentence according to the c' prediction (as measured by embedding distance to a list of pre-defined positive/negative words). This model would seem very reasonable to a human, yet it is non-interpretable for the same reason. To recover a “native” neural model, we could train a selector s to imitate $c' \rightarrow s'$ via teacher-student distillation (Hinton et al., 2015), and the innocent-looking $s \rightarrow c$ rationale model

¹In fact, the extended rationale *helps* disguise the problem by appearing as much more reasonable.

remains equally non-interpretable.

Even without the explicit multi-stage supervision above, a sufficiently flexible selector s (e.g. a neural network) can implicitly learn the $c' \rightarrow s'$ model and essentially control the learning of the classifier c , in which case the bottleneck of succinct rationale affords no benefits of interpretability. So why does interpretability get lost (or fail to emerge)? The issue arises from not understanding the *rationale selection process*, i.e. selector s . If it is well-understood, we could determine its true logic to be $c' \rightarrow s'$ and reject it. Conversely, if we cannot understand *why* a particular rationale is selected, then accepting it (and the resulting prediction) at face value is not really any different from accepting an end-to-end prediction at face value.

In addition, the selector-classifier decomposition suggests that the selector should be an “unbiased evidence collector”, i.e. scanning through the input and highlighting all relevant information, while the classifier should deliberate on the evidence for each class and make the decision. Verifying this role of the selector would again require its interpretability.

Finally, considering the rationale model as a whole, we could also argue that the rationale selector *should* be interpretable. It is already accepted that the classifier can remain a black-box. If the selector is also not interpretable, then exactly what about the model is interpretable?

Architecturally, we can draw an analogy between the rationale in rationale models and the embedding representation in a typical end-to-end classifier produced at the penultimate layer. A rationale is a condensed feature extracted by the selector and used by the classifier, while, for example in image models, the image embedding is the semantic feature produced by the feature extractor and used by the final layer of linear classifier. Furthermore, both of them exhibit some interpretable properties: rationales represent the “essence” of the input, while the image embedding space also seems semantically organized (e.g. Figure 2 showing ImageNet images organized in the embedding space). However, this embedding space is rarely considered on its own as the explanation for a prediction, exactly because the feature extractor is a black-box. Similarly, the rationales by default should not qualify as the explanation either, despite its textual nature.

Finally, from a practical perspective, explanations should help humans understand the model’s input-output behavior. Such a purpose is fulfilled



Figure 2: Embedding space visualization of an ImageNet classifier. Image from <https://cs.stanford.edu/people/karpathy/cnnembed/>.

when the human understands not only why an explanation leads to an output, but also *how* the explanation is generated from the input in the first place. Our emphasis on understanding the rationale selection process fulfills the latter requirement. Such a perspective is also echoed by Pruthi et al. (2020), who argued that the practical utility of explanations depends crucially on human’s capability of understanding how they are generated.

4 Empirical Investigation

As discussed above, truly interpretable rationale models require an understanding of the rationale selection process. However, since the selector is a sequence-to-sequence model, for which there is no standard methods for interpretability, we focus on a “necessary condition” setup of understanding the input-output behavior of the model in our empirical investigation. Specifically, we investigate rationale selection changes in response to meaning-preserving non-adversarial perturbation of individual words in the input sentence.

4.1 Setup

On the 5-way SST dataset (Socher et al., 2013), we trained two rationale models, a continuous relaxation (CR) model (Bastings et al., 2019) and a policy gradient (PG) model (Lei et al., 2016). The PG model directly generates binary (i.e. hard) rationale selection. The CR model uses a $[0, 1]$ continuous value to represent selection and scales the word embedding by this value. Thus, we consider a word being selected as rationale if this value is non-zero. Our CR model achieves 47.3% test accuracy with 24.9% rationale selection rate (i.e. percentage

of words in the input selected as rationale), and PG model 43.3% test accuracy with 23.1% rationale selection rate, consistent with those obtained by Bastings et al. (2019, Figure 4). Additional details are in Appendix A.

4.2 Sentence Perturbation Procedure

The perturbation procedure changes a noun, verb, or adjective as parsed by NLTK² (Loper and Bird, 2002) with two requirements. First, the new sentence should be natural (e.g., “I *observed* a movie” is not). Second, its meaning should not change (e.g. adjectives should not be replaced by antonyms).

For the first requirement, we [MASK] the candidate word and use the pre-trained BERT (Devlin et al., 2019) to propose 30 new choices. For the second requirement, we compute the union of words in the WordNet synset associated with each definition of the candidate words (Fellbaum, 1998). If the two sets share no common words, we mark the candidate invalid. Otherwise, we choose the top BERT-predicted word as the replacement.

We run this procedure on the SST test set, and construct the perturbed dataset from all valid replacements of each sentence. Table 1 lists some example perturbations (more in Appendix B). Table 2 shows the label prediction distribution on the original test set along with changes due to perturbation in parentheses, and confirms that the change is overall very small. Finally, a human evaluation checks the perturbation quality, detailed in Appendix C. For 100 perturbations, 91 were rated to have the same sentiment value. Furthermore, on all 91 sentences, the same rationale is considered adequate to support the prediction after perturbation as well.

	A pleurably jacked-up piece/slice of action moviemaking .
	The use/usage of CGI and digital ink-and-paint make the thing look really slick .

Table 1: Sentence perturbation examples, with the original word in **bold** replaced by the word in *italics*.

	0	1	2	3	4
CR	8.0 (-0.6)	41.5 (+1.5)	8.9 (-0.4)	28.6 (+1.0)	13.0 (-1.5)
PG	8.6 (-1.6)	40.7 (-1.3)	1.6 (-0.2)	33.9 (+5.0)	15.2 (-1.9)

Table 2: The percentage of predicted labels on the original test set, as well as the differences to the that on the perturbation sentences in parentheses.

4.3 Results

Now we study the effects of perturbation on rationale selection change (i.e. an originally selected

²i.e. NN, NNS, VB, VBG, VBD, VBN, VBP, VBZ, and JJ

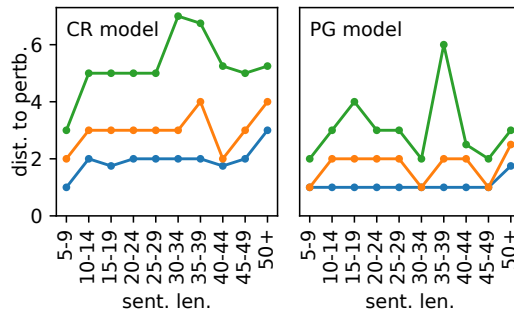


Figure 3: Scatter plots showing three quartiles of distance between indirect rationale change to perturbation, grouped by sentence length.

word getting unselected or vice versa). We use only perturbations that maintain the model prediction, as in this case, the model is expected to use the same rationales according to human evaluation.

Qualitative Examples Table 3 shows examples of rationale changes under perturbation (more in Appendix D). Indeed, minor changes can induce nontrivial rationale change, sometimes far away from the perturbation location. Moreover, there is no clear relationship between the words with selection change and the perturbed word.

PG	The story/narrative loses its bite in a last-minute happy ending that 's even less plausible than the rest of the picture .
PG	A pleasant ramble through the sort of idiosyncratic terrain that Errol Morris has/have often dealt with ... it does possess a loose , lackadaisical charm .
CR	I love the way that it took chances and really asks you to take these great/big leaps of faith and pays off .
CR	Legendary Irish writer/author Brendan Behan 's memoir , Borstal Boy , has been given a loving screen transferral .

Table 3: Rationale change example. Words selected in the original only, perturbed only, and both are shown in red, blue, and green, respectively.

Rationale Change Freq. Quantitatively, we first study how often rationales change. Table 4 shows the count frequency of selection changes. Around 30% (non-adversarial) perturbations result in rationale change (i.e. non-zero number of changes). Despite better accuracy, the CR model is less stable and calls for more investigation into its selector.

# Change	0	1	2	3	4	≥ 5
CR	66.5%	25.5%	6.8%	1.0%	0.1%	0.1%
PG	77.4%	21.4%	1.1%	0.1%	0%	0%

Table 4: Frequency of number of selection changes.

Locations of Selection Change Where do these changes occur? 29.6% and 78.3% of them happen at the perturbed word for the CR and PG models respectively. For the CR model, over 70% of rationale changes are due to replacements of *other* words; this statistic is especially alarming. For

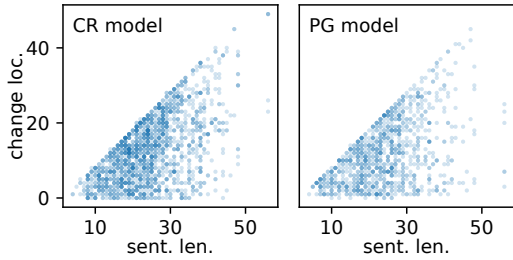


Figure 4: Locations of all selection changes, with each one shown as a dot.

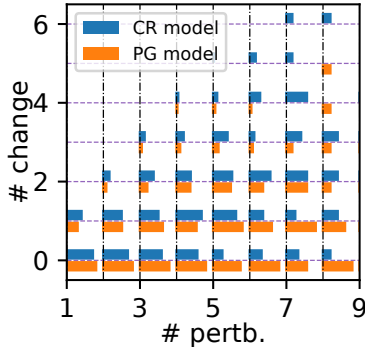


Figure 5: For sentences with a certain number of valid perturbations, the corresponding column of bar chart shows the count frequency of perturbations that result in any rationale change.

these indirect changes, Figure 3 shows the quartiles of distances to the perturbation for varying sentence lengths. They are relatively constant throughout, suggesting that the selection uses mostly local information. However, the “locality size” for CR is about twice as large, and changes often occur five or more words away from the perturbation.

We also compute the (absolute) location of the rationale changes, as plotted in Figure 4, where each dot represents an instance. The rationale changes are distributed pretty evenly in the sentence, making it hard to associate particular perturbation properties to the resulting selection change location.

Sentence-Level Stability Are all the rationale changes concentrated on a few sentences for which every perturbation is likely to result in a change, or are they spread out across many sentences? We measure the *stability* of a sentence by the number of perturbations inducing rationale changes. Obviously, a sentence with more valid perturbations is likely to also have more change-inducing ones, so we plot the frequency of sentences with a certain stability value separately for different total numbers of perturbations in Figure 5. There are very few highly unstable sentences, suggesting that the selection change is a common phenomenon to most of the sentences, further adding to the difficulty of

a comprehensive understanding of the selector.

Part of Speech Analysis Our final automated analysis studies the part-of-speech (POS) composition of selection changes. As Table 5 shows, adjectives and adverbs are relatively stable, as expected because they encode most sentiments. By contrast, nouns and verbs are less stable, probably because they typically represent factual “content” that is less important for prediction. The CR model is especially unstable for other POS types such as determiner and preposition. Overall, the instability adds to the selector complexity and could even function as subtle “cues” described in Section 3.

User Study on Selector Understanding While the automated analyses reveal potential obstacles to selector understanding, ultimately the problem is the lack of understanding by users. The most popular way to understand a model is via input-output examples (Ribeiro et al., 2020; Booth et al., 2021), and we conduct a user study in which we ask participants (grad students with ML knowledge) to match rationale patterns with sentences before and after perturbation on 20 instances, after observing 10 true model decisions (details in Appendix E). Unsurprisingly, participants get 45 correct out of 80 pairs, basically at the random guess level, even as some participants use reasons related to grammar and atypical word usage (which are apparently ineffective), along with “lots of guessing”. This result confirms the lack of selector understanding even under minimal perturbation, indicating more severity for completely novel inputs.

5 Conclusion

We argue against the commonly held belief that rationale models are inherently interpretable by design. We present several reasons, including a counter-example showing that a reasonable-looking model could be as non-interpretable as a black-box. These reasons imply that the missing piece is an understanding of the *rationale selection process* (i.e. the selector). We also conduct a (non-adversarial) perturbation-based study to investigate the selector of two rationale models, in which automated analyses and a user study confirm that they are indeed hard to understand. In particular, the higher-accuracy model (CR) fares worse in most aspects, possibly hinting at the performance-interpretability trade-off (Gunning and Aha, 2019). These results point to a need for more rigorous analysis of interpretability in neural rationale models.

POS (frequency)	noun (19.2%)	verb (14.3%)	adj. (10.1%)	adv. (5.8%)	proper n. (4.4%)	pron. (4.9%)	other (41.3%)
CR change / all	37.1% / 34.3%	21.9% / 16.0%	14.2% / 24.8%	8.9% / 11.3%	3.5% / 5.8%	2.5% / 1.0%	11.9% / 6.8%
PG change / all	42.7% / 33.6%	30.2% / 16.6%	20.6% / 30.6%	2.4% / 12.9%	1.8% / 3.4%	0.4% / 0.5%	1.9% / 2.4%

Table 5: Part of speech (POS) statistics. The top row shows the POS composition of the test set sentences. The bottom two rows show POS composition for changed rationale words and for all rationale words.

References

- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. In *Advances in neural information processing systems*, volume 31, pages 9505–9515.
- Jasmijn Bastings, Wilker Aziz, and Ivan Titov. 2019. [Interpretable neural predictions with differentiable binary variables](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2963–2977, Florence, Italy. Association for Computational Linguistics.
- Serena Booth, Yilun Zhou, Ankit Shah, and Julie Shah. 2021. Bayes-trex: a bayesian sampling approach to model transparency by example. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- 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. 2019. [Eraser: A benchmark to evaluate rationalized nlp models](#).
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. Pathologies of neural models make interpretations difficult. *arXiv preprint arXiv:1804.07781*.
- David Gunning and David Aha. 2019. Darpa’s explainable artificial intelligence (xai) program. *AI Magazine*, 40(2):44–58.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. 2018. [Evaluating feature importance estimates](#). *CoRR*, abs/1806.10758.
- Alon Jacovi and Yoav Goldberg. 2020. [Aligning faithful interpretations with their social attribution](#). *CoRR*, abs/2006.01067.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C Wallace. 2020. Learning to faithfully rationalize by construction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4459–4473.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. [Rationalizing neural predictions](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas. Association for Computational Linguistics.
- Edward Loper and Steven Bird. 2002. [Nltk: The natural language toolkit](#). *CoRR*, cs.CL/0205028.
- Scott Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*.
- Danish Pruthi, Bhuwan Dhingra, Livio Baldini Soares, Michael Collins, Zachary C. Lipton, Graham Neubig, and William W. Cohen. 2020. [Evaluating explanations: How much do explanations from the teacher aid students?](#) *CoRR*, abs/2012.00893.
- 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.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of nlp models with checklist. *arXiv preprint arXiv:2005.04118*.
- 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.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.

- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Mo Yu, Shiyu Chang, Yang Zhang, and Tommi S Jaakkola. 2019. Rethinking cooperative rationalization: Introspective extraction and complement control. *arXiv preprint arXiv:1910.13294*.
- Yilun Zhou, Serena Booth, Marco Tulio Ribeiro, and Julie Shah. 2022a. Do feature attribution methods correctly attribute features? In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*. AAAI.
- Yilun Zhou, Marco Tulio Ribeiro, and Julie Shah. 2022b. [Exsum: From local explanations to model understanding](#).

A Additional Details on the Experimental Setup

A.1 Training

The models we train are as implemented in (Bastings et al., 2019). The hyperparameters we use are 30 percent for the word selection frequency when training the CR model and L_0 penalty weight 0.01505 when training the PG model. Training was done on a MacBook Pro with a 1.4 GHz Quad-Core Intel Core i5 processor and 8 GB 2133 MHz LPDDR3 memory. The training time for each model was around 15 minutes. There are a total of 7305007 parameters in the CR model and 7304706 parameters in the PG model. The hyperparameter for the CR model is the word selection frequency, ranging from 0% to 100%, whereas the hyperparameter for the PG model is the L_0 penalty weight which is a nonnegative real number (for penalizing gaps in selections).

These hyperparameter were configured with the goal that both models would select a similar fraction of total words as rationale. This was done manually. Only one CR model was trained (with the word selection frequency set to 30 percent). Then, a total of 7 PG models were trained, with L_0 penalty weight ranging from 0.01 to 0.025. Then, the closest matching result to the CR model in terms of word selection fraction, which was an L_0 penalty of 0.01505, was used.

The CR model (with 30% word selection frequency) achieves a 47.3% test accuracy with a 24.9% rationale selection rate, and the PG model (with L_0 penalty of 0.01505) achieves a 43.3% test accuracy with a 23.1% selection rate, consistent with those obtained by Bastings et al. (2019, Figure 4). The CR model achieves a validation accuracy of 46.0% with a 25.1% rationale selection rate, and the PG model achieves a 41.1% validation accuracy with a 22.9% selection rate, comparable to the test results.

A.2 Dataset

We use the Stanford Sentiment Treebank (SST, Socher et al., 2013) dataset with the exact same preprocessing and train/validation/test split as given by Bastings et al. (2019). There are 11855 total entries (each are single sentence movie reviews in English), split into a training size of 8544, a validation size of 1101, and a test size of 2210. The label distribution is 1510 sentences of label 0 (strongly negative), 3140 of label 1 (negative), 2242 of label 2 (neutral), 3111 of label 3 (positive), and 1852 of label 4 (strongly positive). We use this dataset as is, and no further pre-processing is done. The dataset can be downloaded from the code provided by Bastings et al. (2019).

A.3 Sentence Perturbation

The data perturbation was done on the same machine with specs described in Appendix A. This procedure was done once and took around an hour. This perturbation was an automated procedure using the BERT and WordNet synset intersection as a heuristic for word substitutions. As a result, we did not collect any new data which requires human annotation or other work.

B Additional Examples of Sentence Perturbation

Table 6 shows ten randomly sampled perturbations.

There are weird resonances between actor and role / <i>character</i> here , and they 're not exactly flattering .
A loving little / <i>short</i> film of considerable appeal .
The film is really not so much / <i>often</i> bad as bland .
A cockamamie tone poem pitched precipitously between swoony lyricism and violent catastrophe ... the most aggressively nerve-wracking and screamingly neurotic romantic comedy in cinema / <i>film</i> history .
Steve Irwin 's method is Ernest Hemmingway at accelerated speed and volume / <i>mass</i> .
The movie addresses a hungry need for PG-rated , nonthreatening family movies / <i>film</i> , but it does n't go too much further .
... the last time I saw a theater full of people constantly checking their watches / <i>watch</i> was during my SATs .
Obvious politics and rudimentary animation reduce the chances / <i>chance</i> that the appeal of Hey Arnold !
Andy Garcia enjoys one of his richest roles in years and Mick Jagger gives his best movie / <i>film</i> performance since , well , Performance .
Beyond a handful of mildly amusing lines ... there just is / <i>be</i> n't much to laugh at .

Table 6: Ten randomly sampled sentence perturbation examples given in a user study, with the original word shown in **bold** replaced by the word in *italics*.

C Description of the Human Evaluation of Data Perturbation

We recruited five graduate students with ML experience (but no particular experience with interpretable ML or NLP), and each participant was asked to answer questions for 20 sentence perturbations, for a total of 100 perturbations. An example question is shown below:

The original sentence (a) and the perturbed sentence (b), as well as the selected rationale on the original sentence (in bold) are:

- There **are weird resonances** between actor and role here , and they 're not **exactly flattering** .
- There are weird resonances between actor and character here , and they 're not exactly flattering .

The original prediction is: negative.

- Should the prediction change, and if so, in which way:
- If yes:
 - Does the changed word need to be included or removed from the rationale?
 - Please highlight the new rationale in red directly on the new sentence.

The study takes less than 15 minutes, is conducted during normal working hours with participants being grad students on regular stipends, and is uncompensated.

D Additional Rationale Change Examples

Table 7 shows additional rationale change examples.

PG	This delicately observed story/tale , deeply felt and masterfully stylized , is a triumph for its maverick director.
PG	Biggie and Tupac is so single-mindedly daring , it puts/put far more polished documentaries to shame.
PG	Somewhere short of Tremors on the modern B-scene : neither as funny nor as clever , though an agreeably unpretentious way to spend/pass ninety minutes .
PG	The film overcomes the regular minefield of coming-of-age cliches with potent/strong doses of honesty and sensitivity .
PG	As expected , Sayles ' smart wordplay and clever plot contrivances are as sharp as ever , though they may be overshadowed by some strong/solid performances .
CR	The animated subplot keenly depicts the inner struggles/conflict of our adolescent heroes - insecure , uncontrolled , and intense .
CR	Funny and , at times , poignant , the film from director George Hickenlooper all takes/take place in Pasadena , “ a city where people still read . ”
CR	It would be hard to think of a recent movie that has/have worked this hard to achieve this little fun.
CR	This road movie gives/give you emotional whiplash , and you 'll be glad you went along for the ride .
CR	If nothing else , this movie introduces a promising , unusual kind/form of psychological horror .

Table 7: Additional rationale change example. Words selected in the original only, perturbed only, and both are shown in **red**, **blue**, and **green**, respectively.

E Description of the User Study on Rationale Change

Participants were first given 10 examples of rationale selections (shown in bold) on the original and perturbed sentence pair made by the model, with one shown below:

orig: **Escapism** in its **purest form** .
 pert: **Escapism** in its **purest kind** .

Then, they were presented with 20 test questions, where each question had two rationale assignments, one correct and one mismatched, and they were asked to determine which was the correct rationale assignment. An example is shown below:

a orig: **Benefits** from a **strong performance** from Zhao , but it 's Dong Jie 's **face** you **remember** at the end .
 pert: Benefits from a **solid** performance from Zhao , but it 's Dong Jie 's **face** you **remember** at the end
 b orig: Benefits from a **strong** performance from Zhao , but it 's Dong Jie 's **face** you **remember** at the end .
 pert: **Benefits** from a **solid performance** from Zhao , but it 's Dong Jie 's **face** you **remember** at the end
 In your opinion, which pair (a or b) shows the actual rationale selection by the model?

In the end, we ask the participants the following question for any additional feedback.

Please briefly describe how you made the decisions (which could include guessing), and your impression of the model's behavior.

The study takes less than 15 minutes, is conducted during normal working hours with participants being grad students on regular stipends, and is uncompensated.