

Evaluating the Faithfulness of Importance Measures in NLP by Recursively Masking Allegedly Important Tokens and Retraining

Andreas Madsen^{1,2} Nicholas Meade^{1,3,*} Vaibhav Adlakha^{1,3,*} Siva Reddy^{1,3,4}

¹ Mila – Quebec AI Institute ² Polytechnique Montréal

³ McGill University ⁴ Facebook CIFAR AI Chair

{firstname.lastname}@mila.quebec

Abstract

To explain NLP models a popular approach is to use importance measures, such as attention, which inform input tokens are important for making a prediction. However, an open question is how well these explanations accurately reflect a model’s logic, a property called *faithfulness*.

To answer this question, we propose Recursive ROAR, a new faithfulness metric. This works by recursively masking allegedly important tokens and then retraining the model. The principle is that this should result in worse model performance compared to masking random tokens. The result is a performance curve given a masking-ratio. Furthermore, we propose a summarizing metric using relative area-between-curves (RACU), which allows for easy comparison across papers, models, and tasks.

We evaluate 4 different importance measures on 8 different datasets, using both LSTM-attention models and RoBERTa models. We find that the faithfulness of importance measures is both model-dependent and task-dependent. This conclusion contradicts previous evaluations in both computer vision and faithfulness of attention literature.

1 Introduction

The ability to explain neural networks benefits both accountability and ethics when deploying models (Doshi-Velez et al., 2017) and helps develop a scientific understanding of what models do (Doshi-Velez and Kim, 2017). Particularly, in NLP, attention (Bahdanau et al., 2015) is often used as an explanation to provide insight into the logical process of a model (Belinkov and Glass, 2019).

Attention, among other methods such as *gradient* (Baehrens et al., 2010; Li et al., 2016) and *integrated gradient* (Sundararajan et al., 2017; Mudrakarta et al., 2018), explain which input tokens

are relevant for a given prediction. This type of explanation is called an importance measure.

A major challenge in the field of interpretability is ensuring that an explanation is *faithful*: “a faithful interpretation is one that accurately represents the reasoning process behind the model’s prediction” (Jacovi and Goldberg, 2020). Unfortunately, importance measures that are claimed to have strong theoretical foundations and are widely used in practice (Bhatt et al., 2019) often later turn out to be questionable (Hooker et al., 2019; Kindermans et al., 2019; Adebayo et al., 2018; Jain and Wallace, 2019; Wiegrefe and Pinter, 2019).

Accurately measuring if an explanation is faithful is therefore paramount. Such *faithfulness* metrics are difficult to develop as the models are too complex to know what the correct explanation is. Doshi-Velez and Kim (2017) says a *faithfulness* metric should use “some formal definition of interpretability as a proxy for explanation quality.”

In this work, we use the definition of *faithfulness* by Samek et al. (2017) and Hooker et al. (2019): if information (input tokens) is truly important, then removing it should result in a worse model performance compared to removing random information (tokens). We build upon the ROAR metric by Hooker et al. (2019), which adds that it is necessary to retrain the model after information is removed, to avoid out-of-distribution issues. Finally, the model performance is compared with removing random information.

A limitation of ROAR is that it is theoretically impossible to measure the faithfulness of an importance measure when dataset redundancies exist. For example, if two tokens are equally relevant but only one of them is identified as important, ROAR fails to remove the second token.

We propose *Recursive ROAR* which solves this limitation. In addition to the *Recursive ROAR* metric, we introduce a summarizing metric (RACU) which aggregates the results into a scalar metric.

*Equal contribution.

We hope that such a metric will make it more feasible to compare importance measures across papers.

Using the proposed faithfulness metrics, we perform a comprehensive comparative study of 4 different importance measures and two popular architectures: BiLSTM-Attention and RoBERTa (Liu et al., 2019). We use 8 different datasets which are commonly used in the faithfulness of *attention* literature (Jain and Wallace, 2019).

Our comparative study reveals that no importance measure is consistently better than others. Instead, we find that faithfulness is both task and model dependent. This is valuable knowledge, as although each importance measure might be equal in faithfulness, they are not equal in computational requirements or understandability to humans.

In particular, we find that *attention* generally provides more sparse explanations than *gradient* or *integrated gradient*. Although their faithfulness may be the same, a sparser explanation is often easier for humans to understand (Miller, 2019).

Computationally speaking, *integrated gradient* is approximately 50 times more expensive than the *gradient* method. This additional complexity is usually justified by being considered more faithful than *gradient*. However, our results indicate that this is rarely a worthwhile trade-off.

2 Related Work

Much recent work in NLP has been devoted to investigating the faithfulness of importance measures, particularly *attention*. In this section, we categorize these faithfulness metrics according to their underlying principle and discuss their drawbacks. ROAR (Hooker et al., 2019) and our Recursive ROAR metrics differ significantly from these approaches.

The works on *attention* are all based on the BiLSTM-Attention models and datasets from Jain and Wallace (2019), they are therefore highly comparable. We use the same models and datasets, while also analyzing RoBERTa.

2.1 Comparing with alternative importance measures

The idea is to compare *attention* with an alternative importance measure, such as *gradient*. The claim is, if there is a correlation this would validate *attention's* faithfulness. Jain and Wallace (2019) specifically compare with the *gradient* method and the *leave-one-out* method. Meister et al. (2021) repeat this experiment in a broader context.

Both Jain and Wallace (2019) and Meister et al. (2021) find that there is little correlation between importance measures and interpret this as attention being not faithful.

Jain and Wallace (2019) does acknowledge the limitations of this approach, as the alternative importance measures are not themselves guaranteed to be faithful. A correlation, or lack of correlation, does therefore not inform about faithfulness. A criticism that we agree with and highlight here.

2.2 Mutate attention to deceive

Jain and Wallace (2019) propose that if there exist alternative attention weights that produce the same prediction, *attention* is unfaithful.

They implement this idea by directly mutating the attention such that there is no prediction change but a large change in *attention* and find that alternative attention distributions exist. Vashishth et al. (2019) and Meister et al. (2021) apply a similar method and achieve similar results.

Wiegrefe and Pinter (2019) find this analysis problematic because the attention distribution is changed directly, thereby creating an out-of-distribution issue. This means that the new attention distribution may be impossible to obtain naturally from just changing the input, and it therefore says little about the faithfulness of attention.

2.3 Optimize model to deceive

Because the *mutate attention to deceive* approach has been criticized for using direct mutation, an alternative idea is to learn an adversarial *attention*.

Wiegrefe and Pinter (2019) investigate maximizing the KL-divergence between normal attention and adversarial attention while minimizing the prediction difference between the two models. By varying the allowed prediction difference, they show that it is not possible to significantly change the attention weights without affecting performance. Importantly, Wiegrefe and Pinter (2019) only use this experiment to invalidate the *mutate attention to deceive* experiments, not to measure faithfulness. However, (Meister et al., 2021) do use this experiment setup as a faithfulness metric.

Pruthi et al. (2020) perform a similar analysis but report a contradictory finding. They find it is possible to significantly change the attention weights without affecting performance. They use this to show that attention is not faithful.

We find this approach problematic because by changing the optimization criteria the analysis is no longer about the standard BiLSTM-attention model (Jain and Wallace, 2019), which is the subject of interest. Therefore, this analysis only works as a criticism of the *mutate attention to deceive* approach, not as an evaluation of faithfulness.

2.4 Known explanations in synthetic tasks

Arras et al. (2022) constructs a purely synthetic task, where the true explanation is known. Evaluating importance measures against this true explanation serves as the faithfulness metric. Unfortunately, this approach cannot be used on real datasets and assumes a well behaved model.

Bastings et al. (2021), a concurrent work to ours, therefore introduce spurious correlations into real datasets, creating partially synthetic tasks. They then evaluate if importance measures can detect these correlations. They conclude, similar to us, that faithfulness is both model and task-dependent.

We believe that this approach is the most valid among the mentioned metrics in the section. However, model behavior, and thereby the explanation behavior, can be drastically different on observations with spurious correlations from those without. This method is therefore limited in scope as it can only evaluate if the importance measure can be used to detect known spurious correlations.

3 ROAR: RemOve And Retrain

To address the shortcomings of the current faithfulness measures as described in Section 2, we base our metric on ROAR (Hooker et al., 2019).

ROAR has been used in computer vision to evaluate the faithfulness of importance measures and to a limited extent in NLP (Pham et al., 2021). The central idea is that if information is truly important, then removing it from the dataset and retraining a model on this reduced dataset should worsen model performance. This can then be compared with an uninformative baseline, where information is removed randomly.

For example, at a step size of 10%, one can remove the top- $\{10\%, 20\%, \dots 90\%$ allegedly important tokens, evaluate the model performance, and compare this with removing $\{10\%, 20\%, \dots 90\%$ random tokens. If the importance measures is faithful, the former should result in a worse model performance than the latter.

This section covers how ROAR is adapted to an

NLP context. Furthermore, we explain the dataset redundancy issue which is solved by our proposed Recursive ROAR metric. Finally, we show that Recursive ROAR is an improvement on ROAR using a synthetic task.

3.1 Adaptation to NLP

ROAR was originally proposed as a faithfulness metric in computer vision. In this context, pixels measured to be important are “removed” by replacing them with an uninformative value, such as a gray pixel (Hooker et al., 2019).

In this work, ROAR is applied to sequence classification tasks. Because these models use tokens, the uninformative value is a special [MASK] token (example in Figure 1). We choose a [MASK] token rather than removing the token to keep the sequence length, which is an information source unrelated to importance measures.

0% The movie is great . I really liked it .
 10% The movie is [MASK] . I really liked it .
 20% The [MASK] is [MASK] . I really liked it .

Figure 1: Example of **ROAR**. The first sentence shows the importance of various tokens. The next two sentences demonstrate the proportion of important tokens replaced by [MASK]. Note, the second sentence is enough to infer the sentiment.

3.2 Recursive ROAR

With ROAR there are two conclusions, either 1) the importance measure is to some degree faithful or, 2) the faithfulness is unknown. The former is observed when the model’s performance is statistically significantly below the random baseline. In the latter case, Hooker et al. (2019) explain that the importance measure can either be not faithful or there can be a dataset redundancy. Recursive ROAR solves this redundancy issue and thereby provides a more informative conclusion.

A dataset redundancy affects the conclusion because the model does not need to use the redundant information. A faithful importance measure would therefore not highlight redundancies as important. After the important information which the importance measure did highlight is removed and the model is retrained, the redundant information can still keep the model’s performance high. An example of this issue is demonstrated in Figure 1.

We solve this issue by recursively recomputing

0%	The	movie	is	great	.	I	really	liked	it	.
10%	The	movie	is	[MASK]	.	I	really	liked	it	.
20%	The	movie	is	[MASK]	.	I	really	[MASK]	it	.

Figure 2: Example of how a redundancy can be removed in **Recursive ROAR** by reevaluating the importance measure. Compare this to Figure 1, where redundancies are not removed and the performance can remain the same, even when the importance measure is faithful.

the importance measure at each iteration of information removal. This way, if the importance measure is faithful, it would quickly mark the redundant information as important after which it would be removed. Note that already masked tokens are kept masked. We call this Recursive ROAR and provide an example in Figure 2.

Note, Recursive ROAR might not remove all redundancies unless the step size is one token. However, because ROAR requires retraining the model, for every evaluation step, this is infeasible. Instead, we approximate it by removing a relative number of tokens. We discuss this more in Appendix F.

3.3 Validation on a synthetic problem

To show that Recursive ROAR provides an optimal faithfulness metric, we validate it on the same generated synthetic problem (with input \mathbf{x} and output y) presented in the original ROAR paper (Hooker et al., 2019):

$$\mathbf{x} = \frac{\mathbf{a}z}{10} + \mathbf{d}\eta + \frac{\epsilon}{10}, \quad y = \begin{cases} 1 & z > 0 \\ 0 & z \leq 0 \end{cases}. \quad (1)$$

Quoting Hooker et al. (2019) “All random variables were sampled from a standard normal distribution. The vectors \mathbf{a} and \mathbf{d} are 16 dimensional vectors that were sampled once to generate the dataset. In \mathbf{a} only the first 4 values have nonzero values to ensure that there are exactly 4 informative features. The values z , η , and ϵ are sampled independently for each example.”

The ground truth removal order is to remove the first 4 features (the specific order does not matter) followed by the remaining irrelevant features. Note that these first 4 features are mutually redundant.

In Hooker et al. (2019), they do not use a specific importance measure. Instead, they use predefined removal orders. This avoids the redundancy issue in the synthetic task, although they do mention it as a limitation. Instead, we use the weights of a linear

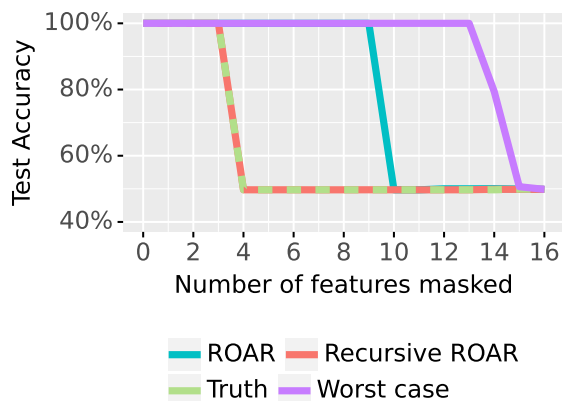


Figure 3: Using the weights of a linear model as the explanation, ROAR and Recursive ROAR are applied to the problem described in (1). In addition, the ground truth and worst case are shown. Recursive ROAR and the ground truth are identical.

model as the importance measure and apply ROAR and Recursive ROAR using this explanation.

Figure 3 shows that Recursive ROAR is identical to the ground truth, while ROAR is worse.

4 Importance Measures

In this section, we describe the importance measures that will be evaluated. We choose these explanations as they are common and computationally feasible to evaluate on every observation.

As *attention* does not attend to the begin-of-sequence token, end-of-sequence token, and auxiliary sequence in paired-sequence problems, these tokens are also not considered for other importance measures. This is to ensure a fair comparison.

Attention These are the attention weights of a BiLSTM-Attention model. We repeat the definitions in Appendix C.1.

While we also look at a transformer-based model which also have internal attention mechanisms, these models do not provide one specific way to convert attention scores into an importance measure. There are proposals to turn the many attention heads into an importance measure (Abnar and Zuidema, 2020). However, these are computationally expensive and requires knowing which layer to select. Performing this analysis is a standalone research topic which we will not answer.

Gradient Let the logits be denoted as $f(\mathbf{x})$. Then the gradient explanation is $\nabla_{\mathbf{x}}f(\mathbf{x})$, where \mathbf{x} is a one-hot-encoding of the input (Baehrens et al.,

2010; Li et al., 2016). To reduce away the vocabulary dimension, we use an L_2 -norm.

Input times Gradient This explanation is $\mathbf{x} \odot \nabla_{\mathbf{x}} f(\mathbf{x})$. Note that because \mathbf{x} is a one-hot encoding, only one element per token will be non-zero. This non-zero element is considered as the explanation.

Integrated Gradient (IG) Sundararajan et al. (2017) argue this to be more faithful, via axiomatic proofs, compared to previous gradient-based methods. A disadvantage is that it is significantly more computationally intensive as it requires computing k gradients. We use $k = 50$ like the original paper (Sundararajan et al., 2017), and use $\mathbf{b} = \mathbf{0}$ as is done in NLP literature (Mudrakarta et al., 2018):

$$\text{IG}(\mathbf{x}) = (\mathbf{x} - \mathbf{b}) \odot \frac{1}{k} \sum_{i=1}^k \nabla_{\tilde{\mathbf{x}}_i} f(\tilde{\mathbf{x}}_i)_c \quad (2)$$

$$\tilde{\mathbf{x}}_i = \mathbf{b} + \frac{i}{k}(\mathbf{x} - \mathbf{b}).$$

5 Experiments

The datasets, performance metrics, and the BiLSTM-attention model are identical to those used in Jain and Wallace (2019) and most other literature evaluating the faithfulness of *attention*. In addition, we use the RoBERTa-base model with the standard fine-tuning procedure (Liu et al., 2019). Details are in Appendix C¹.

We report model performance on the 8 studied datasets in Table 1. Below, we provide a short description of each dataset. We provide additional details in Appendix B.

1. Two sentiment tasks: SST (Socher et al., 2013) and IMDB (Maas et al., 2011).
2. Two tasks with long-sequences: Diabetes and Anemia (Johnson et al., 2016). These datasets contain many redundancies.
3. A paired-sequence class: SNLI (Bowman et al., 2015).
4. *bAbI* (Weston et al., 2016) task 1 to 3. These are synthetic paired-sequence problems.

5.1 Supporting experiments

In Appendix G, we compare *ROAR* and *Recursive ROAR*. These results show dataset redundancies interfere with *ROAR*. For example, with the Diabetes dataset, only by using *Recursive ROAR* can *gradient* be measured to be faithful.

¹Code is available at <https://github.com/AndreasMadsen/nlp-roar-interpretability>

Dataset	Sequence length	Performance [%]	
		LSTM	RoBERTa
Anemia	2267	88 ^{+1.1} _{-2.2}	86 ^{+0.6} _{-0.7}
Diabetes	2207	81 ^{+2.2} _{-2.9}	76 ^{+0.7} _{-0.6}
IMDB	181	90 ^{+0.4} _{-0.7}	95 ^{+0.2} _{-0.2}
SNLI	16	78 ^{+0.2} _{-0.3}	91 ^{+0.1} _{-0.1}
SST	20	82 ^{+0.6} _{-1.0}	94 ^{+0.3} _{-0.3}
bAbI-1	38	100 ^{+0.0} _{-0.1}	100 ^{+0.0} _{-0.0}
bAbI-2	96	68 ^{+9.1} _{-19.1}	100 ^{+0.1} _{-0.1}
bAbI-3	308	60 ^{+6.5} _{-4.9}	81 ^{+6.8} _{-20.0}

Table 1: Model performance scores and sequence-length for each dataset. Performance is averaged over 5 seeds with a 95% confidence interval. Following Jain and Wallace (2019), we report performance as macro-F1 for SST, IMDB, Anemia and Diabetes, micro-F1 for SNLI, and accuracy for bAbI.

The average sequence-length is for the BiLSTM-attention model, for the RoBERTa model the number will be higher but with inputs truncated at 512 tokens.

In Appendix F, we avoid the approximation of removing a relative number of tokens at 10% increments by instead removing exactly one token in each iteration. These results show that the approximation does affect the results, but not the conclusions that can be drawn from the results.

In Appendix E, we report the sparsity of each importance measure and find that *attention* is significantly more sparse than other importance measures. If the faithfulness is equal, this may make it more desirable as sparse explanations are more understandable to humans (Miller, 2019).

5.2 Main experiment: Recursive ROAR

To evaluate the faithfulness of importance measures, we apply *Recursive ROAR* to all datasets and both models. The results are presented in Figure 4 and discussed in Section 6.

In Appendix D, we report the compute times. Because BiLSTM-Attention is a small model and RoBERTa-base is only fine-tuned, Recursive ROAR is feasible when importance measure can be evaluated on every observation. For some importance measures, like SHAP (Lundberg and Lee, 2017), which have exponential compute complexity, ROAR would not be feasible. Additionally, for large language models, like T5 (Raffel et al., 2020), ROAR would also be difficult to apply as fine-tuning these models is generally challenging.




Figure 4: Recursive ROAR results, showing model performance at $x\%$ of tokens masked. A model performance below *random* indicates faithfulness, while above or similar to *random* indicates a non-faithful importance measure. Performance is averaged over 5 seeds with a 95% confidence interval.

5.2.1 How to interpret

If the model performance of a given importance measure is below the random baseline, then this indicates a faithful importance measure. Note that “faithful” is not absolute, rather we measure the degree of faithfulness. However, if the model performance is not statistical significant below the random baseline, then the importance measure is not considered to be faithful. With the (*Not Recursive*) ROAR measure, this latter case would be

inconclusive as the faithfulness could be hidden by dataset redundancies.

Figure 4 also presents the model performance at 100% masking, which provides a lower bound for the model performance and is helpful as the datasets are often biased. These biases come from unbalanced classes or the secondary sequence for the paired-sequence tasks (Gururangan et al., 2018). For these datasets, sequence-length bias is not a concern Appendix B.3.

6 Important Findings

Based on the results in Figure 4 and Table 2, we highlight the following important findings.

Faithfulness is model-dependent. In particular, the faithfulness with SNLI is highly model-dependent as seen in Table 2. Furthermore, comparing the faithfulness between the two models, the faithfulness of *Gradient* on IMDB and *Integrated Gradient* on bAbI-3 is significantly affected by the model architecture.

Faithfulness is task-dependent. For BiLSTM-Attention, in Table 2, *Attention* is best for SNLI while *Input times Gradient* and *Integrated Gradient* is best for SST.

For RoBERTa, *Integrated Gradient* is best for IMDB and SNLI, while *Gradient* is best for bAbI-1 and bAbI-2. In fact, *Integrated Gradient* is worst in all bAbI tasks.

Attention can be faithful. In Table 2, *Attention* is among the top explanations in terms of faithfulness, except for SST. This contradicts many of the previous results mentioned in Section 2, which found attention to be unfaithful.

Because attention is computationally free and attention is more sparse (Appendix E), which is important for human understanding (Miller, 2019), attention can be an attractive explanation.

Integrated Gradient is not necessarily more faithful than Gradient or Input times Gradient. For BiLSTM-Attention, in Table 2, bAbI-2, bAbI-3, and SNLI has least one gradient-based importance measure which is significantly more faithful than *Integrated Gradient*. For RoBERTa, we find the same for bAbI-2 and bAbI-3. These results contradicts the claim that *Integrated Gradient* is theoretically superior (Sundararajan et al., 2017). This is a valuable finding, as *Integrated Gradient* is significantly more computationally expensive than other gradient-based importance measure.

Importance measures often work best for the top-20% most important tokens. In Figure 4, we observe that the largest drop tends to happen at about 10% or 20% tokens masked. This indicates that importance measures are best at ranking the most important tokens, while for less important tokens, they become noisy. This is particularly observed in bAbI for both models and Diabetes with the BiLSTM-Attention model.

Class leakage can cause the model performance to increase. Because the importance measures explain predictions of the target label, they can leak the target label when allegedly important tokens are masked.

Consider a sentiment classification task. If an importance measure indicates that the word *bad* is a strong indicator of negative sentiment, then in the next iteration *bad* would be masked in negative sentences. This means the presence of *bad* now leaks the true label (positive sentiment) which may increase the performance.

This issue is particularly observed with bAbI-3 using RoBERTa in Figure 4, where the performance increases slightly at 60% tokens masked. This issue affects both ROAR and Recursive ROAR (Appendix G). In fact, it likely affects most faithfulness metrics. However, Recursive ROAR can mitigate this issue to some extent. We discuss this more in Appendix A.

7 Conclusion

We show that Recursive ROAR is an improvement on ROAR. In a synthetic setting, Recursive ROAR matches the ground truth, while ROAR does not. Additionally, we argue why other faithfulness metrics may be either invalid or limited in scope.

We then use Recursive ROAR to measure the faithfulness of the most common importance measures, including attention. This is done on both recurrent and transformer-based neural models. In general, we find that the faithfulness of importance measures is both model-dependent and task-dependent. This means that no general recommendation can be made for NLP practitioners considering the current importance measures. Instead, it is necessary to measure the faithfulness of different importance measures given a task and a model.

Because Recursive ROAR works on real-world datasets and not just synthetic problems, we hope it can serve as a standardized benchmark for the faithfulness of importance measures in NLP.

8 Limitations

Recursive ROAR requires the model to be retrained. This means it is not possible to evaluate the faithfulness of a specific model instance, rather we evaluate the faithfulness of the model architecture. The confidence intervals we provide then inform us about what can be statistically expected in terms of the faithfulness for a model instance.

The retraining dependence also means Recursive ROAR can only measure the faithfulness of a task-model combination that is feasible to train/fine-tune repeatedly and importance measures that are feasible to compute across the entire dataset.

A second category of limitation comes from the use of masking. In particular, if the dataset is heavily biased, then the performance at 100% will remain high. This can happen if for example the sequence length is a good predictor of the class. In principle, this means that no tokens are important. Therefore, we can't comment on the faithfulness of an importance measure in that context. In such a case, the faithfulness metric in (3) should become unstable (in theory division by zero, but in practice chaotic values) and result in a large confidence interval.

As discussed in the previous section, because the importance measures explain the target class, they can leak the class information when used to mask input features. This can make an importance measure appear less faithful than it actually is. However, this issue cannot make an importance measure appear more faithful than it is (see Appendix A for more discussion).

Furthermore, while we believe Recursive ROAR provides a useful metric for faithfulness, only measuring faithfulness is not enough for an explanation to be used in production settings (Doshi-Velez and Kim, 2017). In addition to faithfulness, one should also evaluate if the explanation is understandable to humans (known as human-groundedness). This is already being done to some extent but is a complex topic (Sen et al., 2020; Hase and Bansal, 2020; Prasad et al., 2021; González et al., 2021; Schuff et al., 2022; Lertvittayakumjorn and Toni, 2019; Nguyen, 2018).

Finally, Doshi-Velez and Kim (2017) argue that explanations should be tested with the final application in mind. Unfortunately, in deployment settings very little evaluation of any kind is done (Bhatt et al., 2019). However, we hope that this work can help establish a metric for faithfulness.

Impact Statement and Ethics

Interpretability itself is paramount to the ethical deployment of machine learning models. Whether this is to proactively ensure that a model performs predictions that align with human values or to retroactively understand what went wrong in a model's prediction (Doshi-Velez and Kim, 2017;

Doshi-Velez et al., 2017).

Providing misleading explanations can be potentially dangerous, as even wrong explanations can be very convincing. To prevent this we need accurate faithfulness metrics, which this paper hopes to provide. However, history has shown that it is notoriously difficult to develop principled faithfulness metrics (Jain and Wallace, 2019; Kindermans et al., 2019; Adebayo et al., 2018; Hooker et al., 2019).

It is always a possibility that a proposed faithfulness metric is flawed, including the one proposed here. If this is not caught it could lead to more misleading explanations. To prevent this, we try to be extra transparent about the limitations of the proposed faithfulness metric, as described in Section 8. In particular, we also advocate for testing an interpretability method in terms of the human-groundedness and application-groundedness before using it in production (Doshi-Velez and Kim, 2017).

Acknowledgements

SR is supported by the Canada CIFAR AI Chairs program and the NSERC Discovery Grant program. Computing resources were provided by Compute Canada.

References

- Samira Abnar and Willem Zuidema. 2020. [Quantifying Attention Flow in Transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Stroudsburg, PA, USA. Association for Computational Linguistics.
- 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 2018-Decem, pages 9505–9515. Curran Associates, Inc.
- Leila Arras, Ahmed Osman, and Wojciech Samek. 2022. [CLEVR-XAI: A benchmark dataset for the ground truth evaluation of neural network explanations](#). *Information Fusion*, 81:14–40.
- David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus Robert Müller. 2010. [How to explain individual classification decisions](#). *Journal of Machine Learning Research*, 11:1803–1831.
- Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by](#)

- jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–15. International Conference on Learning Representations, ICLR.
- Jasmijn Bastings, Sebastian Ebert, Polina Zablotskaia, Anders Sandholm, and Katja Filippova. 2021. "Will You Find These Shortcuts?" A Protocol for Evaluating the Faithfulness of Input Saliency Methods for Text Classification. *arXiv*.
- Yonatan Belinkov and James Glass. 2019. *Analysis Methods in Neural Language Processing: A Survey*. *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José M. F. Moura, and Peter Eckersley. 2019. *Explainable Machine Learning in Deployment*. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 648–657.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. *A large annotated corpus for learning natural language inference*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Finale Doshi-Velez and Been Kim. 2017. *Towards A Rigorous Science of Interpretable Machine Learning*. *arXiv*.
- Finale Doshi-Velez, Mason Kortz, Ryan Budish, Christopher Bavitz, Samuel J. Gershman, David O’Brien, Stuart Shieber, Jim Waldo, David Weinberger, and Alexandra Wood. 2017. *Accountability of AI Under the Law: The Role of Explanation*. *SSRN Electronic Journal*, Online.
- Ana Valeria González, Anna Rogers, and Anders Søgaard. 2021. *On the Interaction of Belief Bias and Explanations*. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2930–2942, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. *Annotation Artifacts in Natural Language Inference Data*. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 107–112, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Peter Hase and Mohit Bansal. 2020. *Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior?* In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5540–5552, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Jan Kindermans, and Been Kim. 2019. *A benchmark for interpretability methods in deep neural networks*. In *Advances in Neural Information Processing Systems*, volume 32.
- Alon Jacovi and Yoav Goldberg. 2020. *Towards Faithfully Interpretable NLP Systems: How Should We Define and Evaluate Faithfulness?* In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Sarthak Jain and Byron C. Wallace. 2019. *Attention is not Explanation*. In *Proceedings of the 2019 Conference of the North*, volume 1, pages 3543–3556, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Liwei H. Wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. *MIMIC-III, a freely accessible critical care database*. *Scientific Data*, 3(1):160035.
- Pieter-Jan Kindermans, Sara Hooker, Julius Adedayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. 2019. *The (Un)reliability of Saliency Methods*. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11700 LNCS, pages 267–280. Springer.
- Piyawat Lertvittayakumjorn and Francesca Toni. 2019. *Human-grounded Evaluations of Explanation Methods for Text Classification*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Section 3, pages 5194–5204, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016. *Visualizing and Understanding Neural Models in NLP*. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 681–691, Stroudsburg, PA, USA. 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*.

- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. *7th International Conference on Learning Representations, ICLR 2019*.
- Scott Lundberg and Su-In Lee. 2017. [A Unified Approach to Interpreting Model Predictions](#). In *Advances in Neural Information Processing Systems*, pages 4766–4775.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *ACL-HLT 2011 - Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. 2021. [Is Sparse Attention more Interpretable?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 122–129, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Tim Miller. 2019. [Explanation in artificial intelligence: Insights from the social sciences](#). *Artificial Intelligence*, 267:1–38.
- Pramod K. Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. [Did the model understand the question?](#) In *ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, volume 1, pages 1896–1906.
- 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)*, volume 1, pages 1069–1078, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Thang M. Pham, Trung Bui, Long Mai, and Anh Nguyen. 2021. [Double Trouble: How to not explain a text classifier’s decisions using counterfactuals synthesized by masked language models?](#)
- Grusha Prasad, Yixin Nie, Mohit Bansal, Robin Jia, Douwe Kiela, and Adina Williams. 2021. [To what extent do human explanations of model behavior align with actual model behavior?](#) In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 1–14, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Danish Pruthi, Mansi Gupta, Bhuwan Dhingra, Graham Neubig, and Zachary C. Lipton. 2020. [Learning to Deceive with Attention-Based Explanations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4782–4793, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21:1–67.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. 2018. [On the convergence of Adam and beyond](#). *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, pages 1–23.
- Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus Robert Müller. 2017. [Evaluating the visualization of what a deep neural network has learned](#). *IEEE Transactions on Neural Networks and Learning Systems*, 28(11):2660–2673.
- Hendrik Schuff, Alon Jacovi, Heike Adel, Yoav Goldberg, and Ngoc Thang Vu. 2022. [Human Interpretation of Saliency-based Explanation Over Text](#). In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 611–636, New York, NY, USA. ACM.
- Cansu Sen, Thomas Hartvigsen, Biao Yin, Xiangnan Kong, and Elke Rundensteiner. 2020. [Human Attention Maps for Text Classification: Do Humans and Neural Networks Focus on the Same Words?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4596–4608, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Richard Socher, John Bauer, Christopher D. Manning, and Andrew Y. Ng. 2013. [Parsing with compositional vector grammars](#). In *ACL 2013 - 51st Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, volume 1, pages 455–465. Association for Computational Linguistics.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic attribution for deep networks](#). In *34th International Conference on Machine Learning, ICML 2017*, volume 7, pages 5109–5118.
- Shikhar Vashishth, Shyam Upadhyay, Gaurav Singh Tomar, and Manaal Faruqui. 2019. [Attention Interpretability Across NLP Tasks](#). *arXiv*.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart Van Merriënboer, Armand Joulin, and Tomas Mikolov. 2016. [Towards AI-complete question answering: A set of prerequisite toy tasks](#). *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*.

Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not Explanation](#). *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 11–20.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Stroudsburg, PA, USA. Association for Computational Linguistics.

A Explanation of class leakage

When importance measures are computed, it is the prediction of the gold label that is explained. For example, for the *Gradient* method, it is $\nabla_{\mathbf{x}} f(x)_y$ that is computed, where \mathbf{x} is the input and y is the gold label.

We want an importance measure for the correct label, as removing the tokens that are relevant for making a wrong prediction, would help the performance of the model. If the gold label was not used, the faithfulness results would be affected by the model performance. As faithfulness and model performance should be unrelated, this is not a desired outcome.

This is a general issue with faithfulness metrics due to how importance measures are calculated in benchmark settings. This is an unfortunate gap between the benchmark-setting and the practical setting where the gold label is unknown. Furthermore, it is rarely documented.

In ROAR and Recursive ROAR, this issue is expressed as an increase in the model performance. Intuitively, it should not be possible for the model performance to increase with more information removed compared to less. However, because the importance measures are w.r.t. the gold label, they can leak the gold label which can increase the model performance.

Thought experiment. Consider the SST dataset, a binary sentiment classification task. Let’s say that the `and` token has a spurious correlation with the positive label (there is some truth to this). Although, clearly the `and` token can appear in both negative and positive sentences.

For example, let’s say that just using the `and` token provides a 60% accurate classification of positive labels. An importance measure would therefore highlight the `and` token as being important for the prediction of positive sentiment. Unfortunately, an importance measure might not consider the `and` token equally important for a negative sentiment (could be due to non-linearity). If all `and` tokens are removed from sentences with positive sentiment as the gold label, the existence of an `and` token is now a perfect predictor of negative sentiment. Hence, the model performance will increase (there will still be negative sentiment sentences without `and` tokens).

Assuming a faithful importance measure, in the next iteration of Recursive ROAR the `and` token would now be important for predicting negative sentiment and would be removed. However, this assumption is rarely completely justified, there is also no guarantee that `and` is considered the most important for all observations. Finally, in the case where a relative number of tokens are masked, the removal of other tokens may leak the gold label.

General issue. As mentioned, the need to use the gold label is a general issue that likely² extends beyond ROAR. However, because ROAR presents a more qualitative metric (Figure 4) where a curve can be observed to increase, this issue is more apparent. Had we just presented the summarizing metric (Table 2), as most faithfulness metrics do, the issue would have been hidden.

B Datasets

The datasets used in this work are listed below. All datasets are public works. There have been made no attempts to identify any individuals. The use is consistent with their intended use and all tasks were already established by Jain and Wallace (2019).

The MIMIC-III dataset (Johnson et al., 2016) is an anonymized dataset of health records. To access this a HIPAA certification is required, which the first author has obtained. Additionally, the MIMIC-III data has not been shared with anyone else, including other authors of this paper.

Below, we provide more details on each dataset. In Table 3, we provide dataset statistics.

²We could not find any documentation for which label is used in relevant non-ROAR metrics, and no code has been published.

Dataset	Size			Performance [%]		
	Train	Validation	Test	LSTM by Jain and Wallace (2019)	LSTM	RoBERTa
Anemia	4262	729	1242	92	88 ^{+1.1} _{-2.2}	86 ^{+0.6} _{-0.7}
Diabetes	8066	1573	1729	79	81 ^{+2.2} _{-2.9}	76 ^{+0.7} _{-0.6}
IMDB	17212	4304	4362	78	90 ^{+0.4} _{-0.7}	95 ^{+0.2} _{-0.2}
SNLI	549367	9842	9824	88	78 ^{+0.2} _{-0.3}	91 ^{+0.1} _{-0.1}
SST	6579	848	1776	81	82 ^{+0.6} _{-1.0}	94 ^{+0.3} _{-0.3}
bAbI-1	8500	1500	1000	100	100 ^{+0.0} _{-0.1}	100 ^{+0.0} _{-0.0}
bAbI-2	8500	1500	1000	48	68 ^{+9.1} _{-19.1}	100 ^{+0.1} _{-0.1}
bAbI-3	8500	1500	1000	62	60 ^{+6.5} _{-4.9}	81 ^{+6.8} _{-20.0}

Table 3: Datasets statistics for single-sequence and paired-sequence tasks. Following Jain and Wallace (2019), we use the same BiLSTM-attention model and report performance as macro-F1 for SST, IMDB, Anemia and Diabetes, micro-F1 for SNLI, and accuracy for bAbI.

B.1 Single-sequence tasks

1. *Stanford Sentiment Treebank (SST)* (Socher et al., 2013) – Sentences are classified as positive or negative. The original dataset has 5 classes. Following Jain and Wallace (2019), we label (1,2) as negative, (4,5) as positive, and ignore the neutral sentences.
2. *IMDB Movie Reviews* (Maas et al., 2011) – Movie reviews are classified as positive or negative.
3. *MIMIC (Diabetes)* (Johnson et al., 2016) – Uses health records to detect if a patient has Diabetes.
4. *MIMIC (Chronic vs Acute Anemia)* (Johnson et al., 2016) – Uses health records to detect whether a patient has chronic or acute anemia.

B.2 Paired-sequence tasks

5. *Stanford Natural Language Inference (SNLI)* (Bowman et al., 2015) – Inputs are premise and hypothesis. The hypothesis either entails, contradicts, or is neutral w.r.t. the premise.
6. *bAbI* (Weston et al., 2016) – A set of artificial text for understanding and reasoning. We use the first three tasks, which consist of questions answerable using one, two, and three sentences from a passage, respectively.

B.3 Class bias and sequence-length bias

Because Recursive ROAR masks tokens the sequence-length remains the same. At 100% masking the only information the model has is the sequence-length. To understand the relevance of the sequence-length, we compare the 100% masking model performance with a basic class-majority

classifier. The results in Table 4 show that the sequence-length does not have much relevance. SNLI does show significant difference but this relates it’s the secondary sequence being a very good predictor on its own, not the sequence length (Gururangan et al., 2018).

Dataset	Majority	LSTM	RoBERTa
Anemia	39%	39% ^{+0.0%} _{-0.0%}	41% ^{+0.0%} _{-0.0%}
Diabetes	45%	45% ^{+0.0%} _{-0.0%}	45% ^{+0.0%} _{-0.0%}
IMDB	34%	33% ^{+0.1%} _{-0.4%}	33% ^{+0.1%} _{-0.3%}
SNLI	34%	67% ^{+0.3%} _{-0.3%}	71% ^{+0.1%} _{-0.1%}
SST	33%	33% ^{+0.0%} _{-0.0%}	33% ^{+0.0%} _{-0.0%}
bAbI-1	15%	15% ^{+0.8%} _{-0.6%}	15% ^{+0.0%} _{-0.0%}
bAbI-2	19%	16% ^{+0.3%} _{-0.4%}	17% ^{+0.4%} _{-0.4%}
bAbI-3	19%	20% ^{+0.8%} _{-1.1%}	18% ^{+1.2%} _{-0.9%}

Table 4: Performance of the class-majority classifier and the BiLSTM-Attention and RoBERTa classifier on the 100% masked dataset. Performance is the standard metric for the dataset. Meaning, macro-F1 for SST, IMDB, Anemia and Diabetes, micro-F1 for SNLI, and accuracy for bAbI.

C Models

C.1 BiLSTM-Attention

The BiLSTM-Attention models, hyperparameters, and pre-trained word embeddings are the same as those from Jain and Wallace (2019). We repeat the configuration details in Table 5.

There are two types of models, single-sequence and paired-sequence, however, they are nearly identical. They only differ in how the context vector

\mathbf{b} is computed. In general, we refer to $\mathbf{x} \in \mathbb{R}^{T \times V}$ as the one-hot encoding of the primary input sequence, of length T and vocabulary size V . The logits are then $f(\mathbf{x})$ and the target class is denoted as c .

C.1.1 Single-sequence

A d -dimensional word embedding followed by a bidirectional LSTM (BiLSTM) encoder is used to transform the one-hot encoding into the hidden states $\mathbf{h}_x \in \mathbb{R}^{T \times 2d}$. These hidden states are then aggregated using an additive attention layer $\mathbf{h}_\alpha = \sum_{i=1}^T \alpha_i \mathbf{h}_{x,i}$.

To compute the attention weights α_i for each token:

$$\alpha_i = \frac{\exp(\mathbf{u}_i^\top \mathbf{v})}{\sum_j \exp(\mathbf{u}_j^\top \mathbf{v})}, \quad \mathbf{u}_i = \tanh(\mathbf{W}\mathbf{h}_{x,i} + \mathbf{b}) \quad (4)$$

where \mathbf{W} , \mathbf{b} , \mathbf{v} are model parameters. Finally, the \mathbf{h}_α is passed through a fully-connected layer to obtain the logits $f(\mathbf{x})$.

C.1.2 Paired-sequence

For paired-sequence problems, the two sequences are denoted as $\mathbf{x} \in \mathbb{R}^{T_x \times V}$ and $\mathbf{y} \in \mathbb{R}^{T_y \times V}$. The inputs are then transformed to embeddings using the same embedding matrix, and then transformed using two separate BiLSTM encoders to get the hidden states, \mathbf{h}_x and \mathbf{h}_y . Likewise, they are aggregated using additive attention $\mathbf{h}_\alpha = \sum_{i=1}^{T_x} \alpha_i \mathbf{h}_{x,i}$.

The attention weights α_i are computed as:

$$\alpha_i = \frac{\exp(\mathbf{u}_i^\top \mathbf{v})}{\sum_j \exp(\mathbf{u}_j^\top \mathbf{v})} \quad (5)$$

$$\mathbf{u}_i = \tanh(\mathbf{W}_x \mathbf{h}_{x,i} + \mathbf{W}_y \mathbf{h}_{y,T_2}),$$

where \mathbf{W}_x , \mathbf{W}_y , \mathbf{v} are model parameters. Finally, \mathbf{h}_α is transformed with a dense layer.

C.2 RoBERTa

We use RoBERTa (Liu et al., 2019) as a transformer-based model due to its consistent convergence. Consistent convergence is helpful as ROAR and Recursive ROAR requires the model to be trained many times. We use the RoBERTa-base pre-trained model and only perform fine-tuning. The hyperparameters are those defined used by Liu et al. (2019, Appendix C) on GLUE tasks. We list the hyperparameters in Table 6.

RoBERTa makes use of a beginning-of-sequence [CLS] token, a end-of-sequence [EOS] token, a separation token [SEP] token, and a masking

token [MASK] token. The masking token used during pre-training is the same token that we use for masking allegedly important tokens.

For the single-sequence tasks, we encode as [CLS] ... *sentence* ... [EOS]. For the paired-sequence tasks, we encode as [CLS] ... *main sentence* ... [SEP] ... *auxiliary sentence* ... [EOS]. Note that when computing the importance measures, only the main sentence is considered. This is to be consistent with the BiLSTM-attention model.

D Compute

In this section, we document the compute times and resources used for computing the results. Unfortunately, our compute infrastructure changed during the making of this paper. The BiLSTM-attention results were computed on V100 GPUs while the RoBERTa results were computed on A100 GPUs. The A100 GPU is significantly faster than the V100 GPU, hence the compute times are not comparable across models. We could have recomputed the BiLSTM-attention results, but doing so would be a waste of resources. We report the machine specifications in Table 7.

The compute times are reported in Table 8. All compute was done using 99% hydroelectric energy.

While the totals in Table 8 may be large, in partial situations only one dataset is usually considered. Additionally, the variance in Figure 4 is quite low, making less seeds an option. Finally, the compute time of *integrated gradient* is approximately 2/3 of the total. As discussed in Section 6, this is rarely worth it. Practical settings may want to not consider *integrated gradient* at all for this reason.

Dataset	Variant	Embedding initialization	Embedding size	nb. of parameters	Batch size	Max epochs
Anemia	Single	Word2Vec trained on MIMIC	300	5 352 158	32	8
Diabetes	Single	Word2Vec trained on MIMIC	300	6 138 158	32	8
IMDB	Single	Pretrained FastText	300	4 218 458	32	8
SNLI	Paired	Pretrained Glove (840B)	300	13 601 939	128	25
SST	Single	Pretrained FastText	300	4 603 658	32	8
bAbI-1	Paired	Standard Normal Distribution	50	55 048	50	100
bAbI-2	Paired	Standard Normal Distribution	50	55 048	50	100
bAbI-3	Paired	Standard Normal Distribution	50	55 048	50	100

Table 5: Details on the BiLSTM-attention models’ hyperparameters. Everything is exactly as done by [Jain and Wallace \(2019\)](#). For all datasets, ASMGrad Adam ([Reddi et al., 2018](#)) is used with default hyperparameters ($\lambda = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$) and a weight decay of 10^{-5} .

Dataset	Variant	Max epochs
Anemia	Single	3
Diabetes	Single	3
IMDB	Single	3
SNLI	Paired	3
SST	Single	3
bAbI-1	Paired	8
bAbI-2	Paired	8
bAbI-3	Paired	8

Table 6: Details on the RoBERTa models’ hyperparameters. RoBERTa ([Liu et al., 2019](#)) is fine-tuned using the RoBERTa-base pre-trained model from HuggingFace ([Wolf et al., 2020](#)) (125M parameters). The hyperparameters are those used by [Liu et al. \(2019\)](#) on GLUE tasks ([Liu et al., 2019](#), Appendix C). The optimizer is AdamW ([Loshchilov and Hutter, 2019](#)), the learning rate has linear decay with a warmup ratio of 0.06, and there is a weight decay of 0.01. Additionally, we use a batch size of 16 and a learning rate of $2 \cdot 10^{-5}$.

Dataset	Importance Measure	Walltime [hh:mm]	
		LSTM	RoBERTa
Anemia	Random	00:09	00:03
	Attention	00:09	–
	Gradient	00:11	00:04
	Input times Gradient	00:11	00:04
	Integrated Gradient	00:44	00:27
Diabetes	Random	00:17	00:05
	Attention	00:17	–
	Gradient	00:23	00:07
	Input times Gradient	00:23	00:07
	Integrated Gradient	01:46	01:09
IMDB	Random	00:05	00:08
	Attention	00:05	–
	Gradient	00:05	00:10
	Input times Gradient	00:05	00:10
	Integrated Gradient	00:20	02:10

BiLSTM-attention	
CPU	4 cores, Intel Gold 6148 Skylake @ 2.4 GHz
GPU	1x NVidia V100 SXM2 (16 GB)
Memory	24 GB
RoBERTa	
CPU	6 cores, AMD Milan 7413 @ 2.65 GHz 128M cache L3
GPU	1x NVidia A100 (40 GB)
Memory	24 GB

Table 7: Compute hardware used for each model. Note, the models were computed on a shared user system. Hence, we only report the resources allocated for our jobs.

E Sparsity

In this section, we analyse the sparsity of each importance measure. While none of the importance measures produce an actual importance for any token, they may have most of the importance assigned to just a few tokens.

This analysis serves two purposes, to show that masking a relative number of tokens is justified and to test if any importance measure are more sparse than others.

Masking a relative number of tokens is justified. If the majority of the importance is assigned to

just a few tokens (e.g. 10 tokens have 99% of the total importance scores), then it would make more sense to perform the non-approximate version of Recursive ROAR where exactly one token is masked in each iteration.

In Figure 6, we look at the sparsity considering the top-10 tokens. We find that that the sparsity is not sufficiently high to justify masking exactly one token in each iteration. For completeness, we

include this analysis in Appendix F.

There are cases where masking exactly one token in each iteration could make sense, for example, for *attention* in bAbI. However, as this is a comparative study among several importance measures and datasets, this is not enough.

Attention is more sparse than others importance measures If a particular importance measure is more sparse than others, while having a




Figure 6: Shows the accumulative importance score relative to the total importance score, for the top-k number of tokens. The metric is averaged over 5 seeds with a 95% confidence interval. Note that datasets are not equal in sequence-length, the scores are therefore hard to compare across datasets. Please refer to Table 1 for statistics on the sequence-length.

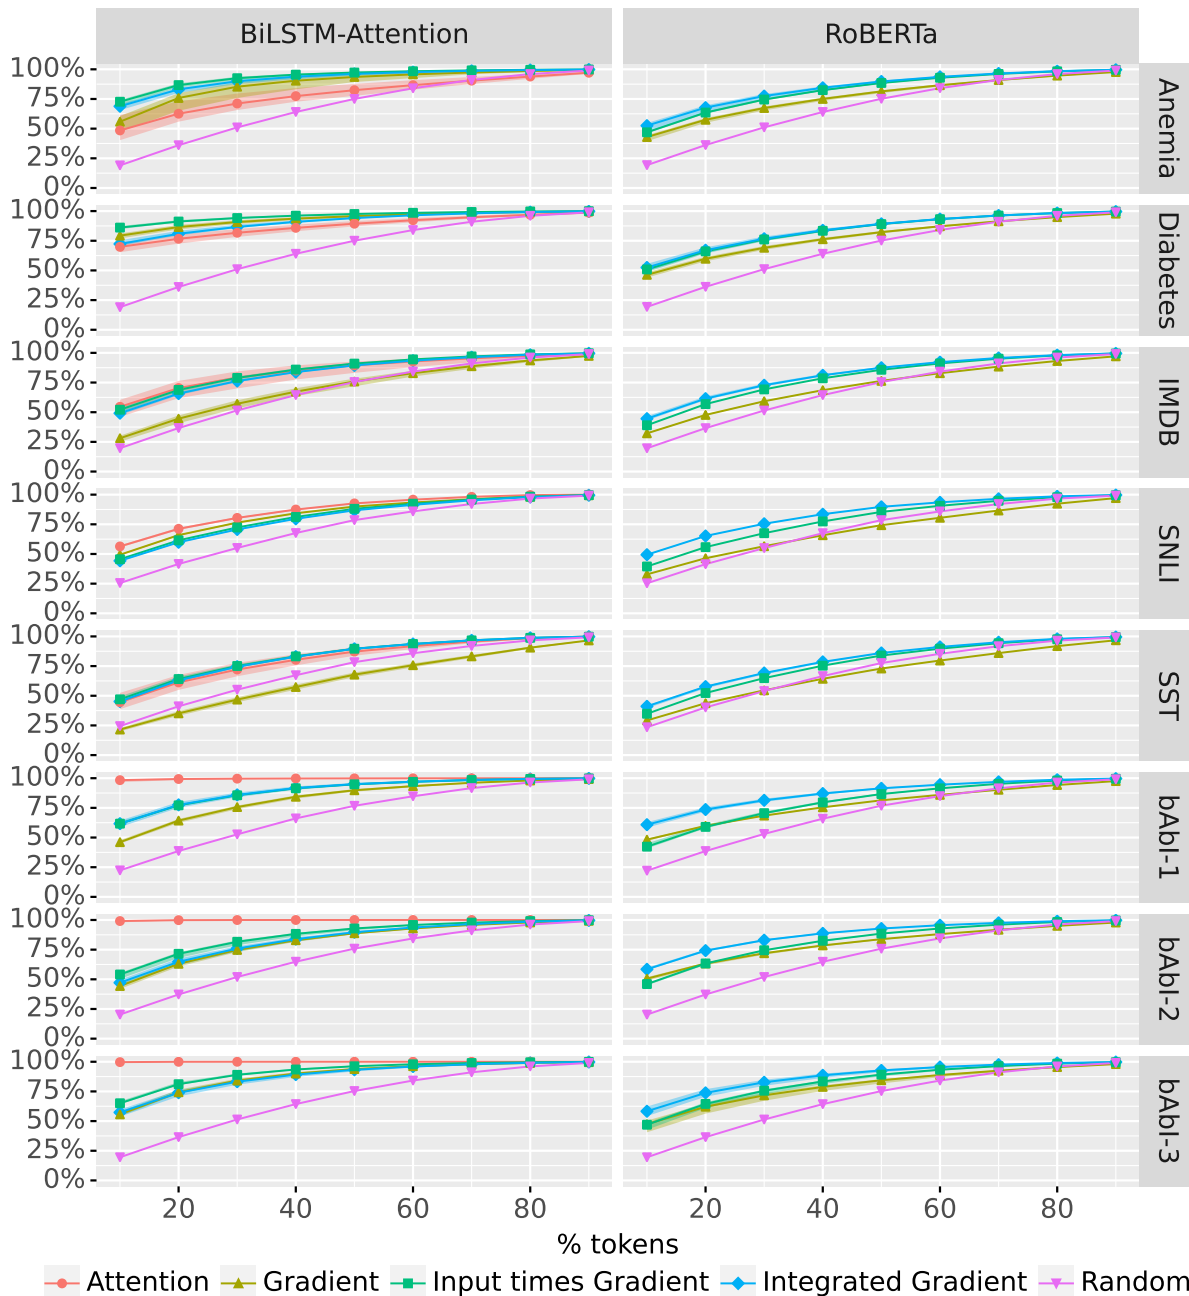


Figure 7: The accumulative importance score relative to the total importance score for the top- $x\%$ number of tokens. The metric is averaged over 5 seeds with a 95% confidence interval.

similar faithfulness, then the more sparse importance measure would be preferable. This is because it is more likely to be understandable to humans (Miller, 2019).

In Figure 7, we look at the sparsity considering a relative number of tokens. We find that for some datasets, in particular bAbI, attention is the most sparse importance measure. Besides this, integrated gradient is usually the most sparse in nearly all cases. However, while the difference in sparsity is often statistically significant we speculate

that the difference is not large enough to cause a difference in practical settings.

F Recursive ROAR with a stepsize of one token

To analyze the effect of masking 10%, as opposed to masking exactly one token in each iteration, we perform the Recursive ROAR experiment with exactly one token masked. The results are in Figure 8. Because this is computationally expensive, we only do this for up to 10 tokens. This

makes it harder to make draw clear conclusions from this experiment, in particular because not all redundancies are removed when only masking 10 tokens.

In general, the results in Figure 8 show that the approximation of masking 10% in each iteration does affect the results. However, we can draw the same conclusions. That being said, some of the conclusions are less obvious because we only look at 10 tokens.

F.1 The results are affected by the approximation

Looking just at RoBERTa, for Diabetes, *Integrated Gradient* yields 65% performance at 10% masking (approximately 51 tokens), while *Integrated Gradient* yields 55% performance at 10 tokens. Similarly for bAbI-3, *Gradient* yields 65% at 10% masking (approximately 30 tokens), while *Gradient* yields 30% at 10 tokens. Both of these cases, shows that a lower performance is achieved earlier when masking one token in each iteration.

This is to be expected, as masking one token in each iteration is more effective for removing redundancies. Were we to complete the experiment to eventually mask all tokens, the faithfulness scores can therefore be expected to be higher.

F.2 The conclusions are the same

In Section 6, we present 5 findings. Here, we briefly show that the same conclusions can be drawn from Figure 8. However, as only 10 tokens are masked they may be less obvious and there may be less evidence.

Faithfulness is model-dependent. Yes, this is most clearly seen for IMDB, where BiLSTM-Attention archives significantly lower performance (higher faithfulness) compared to RoBERTa.

Faithfulness is task-dependent. Yes, looking at BiLSTM-Attention, for IMDB *Integrated Gradient* is the worst importance measure. However, for the bAbI tasks *Integrated Gradient* is among the best importance measures.

Attention can be faithful. Yes, particularly for bAbI, IMDB, and Diabetes attention is faithful.

Integrated Gradient is not necessarily more faithful than Gradient or Input times Gradient. Yes, considering BiLSTM-Attention, IMDB *Integrated Gradient* is significantly worse than other

explanations. For most datasets, *Integrated Gradient* has similar faithfulness as other importance measures.

Importance measures often work best for the top-20% most important tokens. As Figure 8 only shows 10 tokens, which is usually below top-20% this is hard to comment on.

Class leakage can cause the model performance to increase. For RoBERTa, in bAbI-3, the *Integrated Gradient* importance measure can be seen to increase performance after 2 tokens are masked.

G ROAR vs Recursive ROAR

As an ablation study we compare ROAR by Hooker et al. (2019) with our Recursive ROAR. Figure 9 shows the comparison for BiLSTM-Attention and Figure 10 shows the comparison for RoBERTa. Recall that for ROAR by Hooker et al. (2019) it is not possible to say that an importance measure is not faithful.

Some datasets have redundancies which affects ROAR. In particular, we find that Diabetes shows a significant difference comparing ROAR with Recursive ROAR. This is both for BiLSTM-Attention (Figure 9) and RoBERTa (Figure 10). For both models, *Gradient* and *Input times Gradient* becomes faithful with Recursive ROAR. Additionally, for RoBERTa the same is the case for *Integrated Gradient*. This is not surprising, as Diabetes contains incredibly long sequences and contains redundancies.

Also, for IMDB, and to a lesser extent SST, there is a clear difference between BiLSTM-Attention and RoBERTa. This too is not surprising, as sentiment can often be inferred from just a single word. However, there are likely to be many positive or negative words in each observation.

Class leakage affects both ROAR and Recursive ROAR. We observe the class leakage issue for ROAR in SNLI with BiLSTM-Attention and for the bAbI tasks with RoBERTa. We observe the issue for Recursive ROAR in bAbI with BiLSTM-Attention. The fact that the issue mostly exists with bAbI is somewhat encouraging, as the bAbI datasets are synthetic. The class leakage issue appears to affect real datasets less.

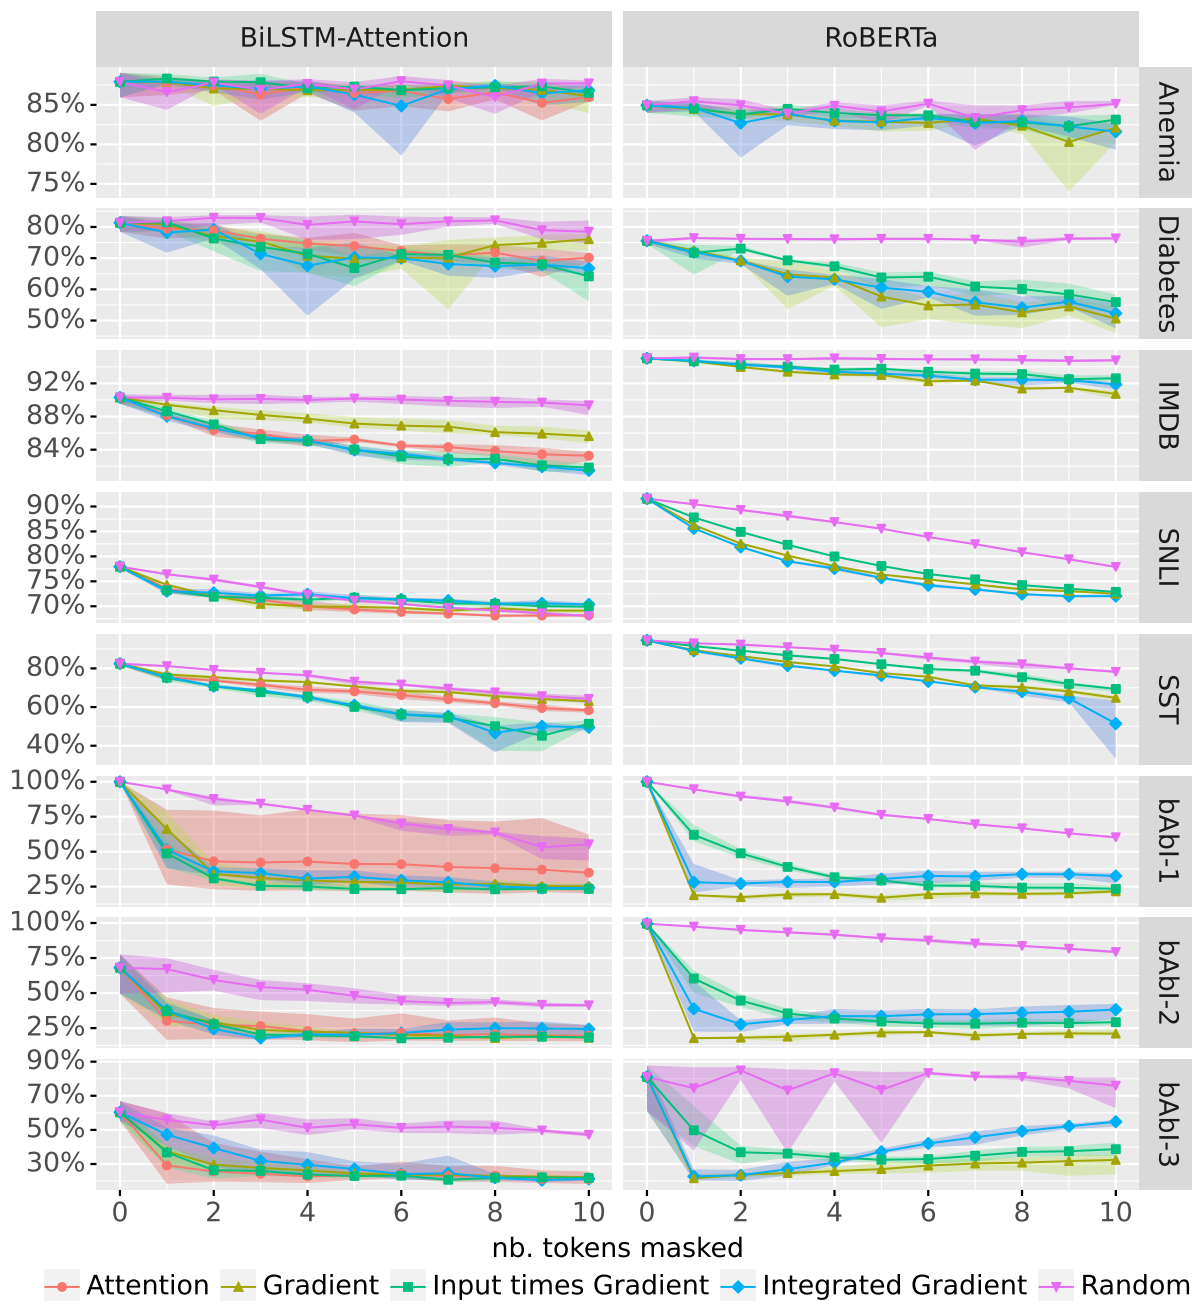


Figure 8: Recursive ROAR results, showing model performance at up to 10 tokens masked. Note that because the datasets have more than 10-tokens, the conclusion one can draw from this plot may change if more tokens were considered. However, in general, a model performance below *random* indicates faithfulness, while above or similar to *random* indicates a non-faithful importance measure. Performance is averaged over 5 seeds with a 95% confidence interval.




Figure 9: ROAR and Recursive ROAR results for **BiLSTM-Attention**, showing model performance at $x\%$ of tokens masked. A model performance below *random* indicates faithfulness. For Recursive ROAR a curve above or similar to *random* indicates a non-faithful importance measure, while for ROAR by Hooker et al. (2019) this case is inconclusive. Performance is averaged over 5 seeds with a 95% confidence interval.

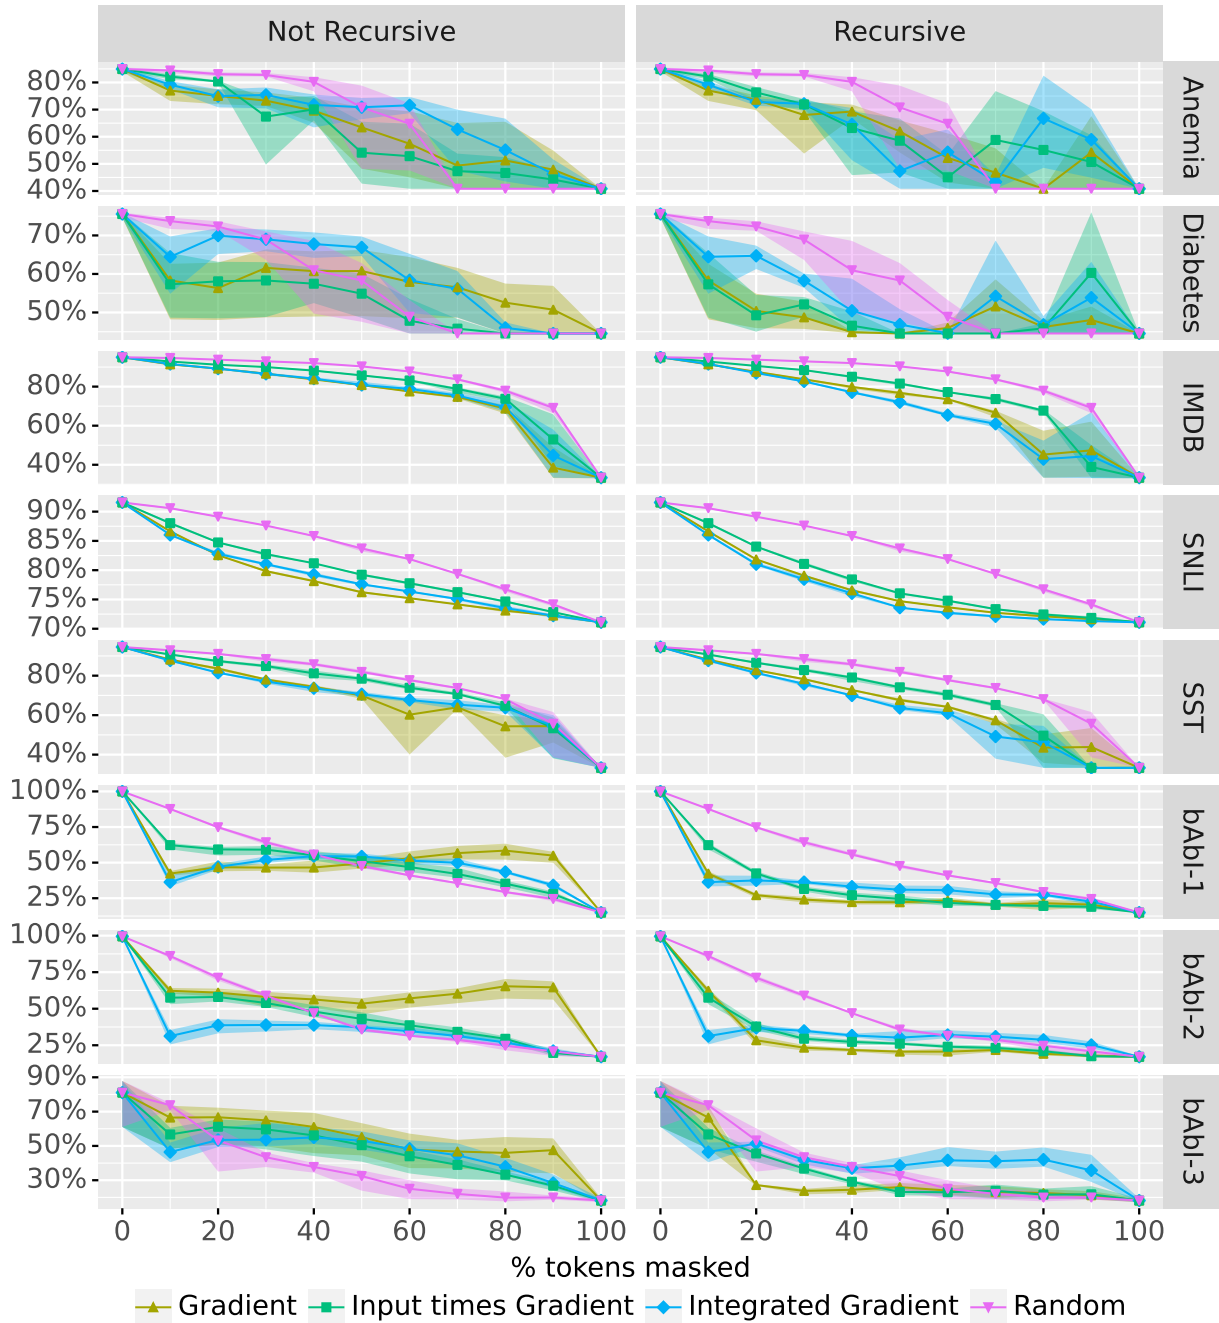


Figure 10: ROAR and Recursive ROAR results for **RoBERTa**, showing model performance at $x\%$ of tokens masked. A model performance below *random* indicates faithfulness. For Recursive ROAR a curve above or similar to *random* indicates a non-faithful importance measure, while for ROAR by Hooker et al. (2019) this case is inconclusive. Performance is averaged over 5 seeds with a 95% confidence interval.