Beyond Fooling: Model Manipulation Under Explanation-aware Training

Tabio Romanski
Radboud University
tabio.romanski@ru.nl

Stjepan Picek
University of Zagreb
Faculty of Electrical Engineering and Computing
& Radboud University
stjepan.picek@ru.nl

Abstract

Feature-level explanations are commonly used to interpret transformer-based NLP models. Still, little is known about how explanation-aware objectives influence model behaviour during training. While prior work has demonstrated training-time manipulation of explanations in vision models, its implications for transformers and token-level explainability remain unexplored. We study training-time manipulation of token-level explanations in transformer-based NLP classifiers and introduce sequence-aware objectives suited to text input. We show that explanation-aware training systematically alters token relevance patterns while largely preserving task accuracy. Importantly, masking and cross-method evaluations reveal that these attribution changes can coincide with shifts in functional model reliance rather than isolated failures of specific explanation methods. Our results suggest that apparent vulnerabilities of feature-level explanations can reflect deeper model adaptations, underscoring the need to consider learning dynamics when interpreting explanation robustness. Our code is available at https://anonymous.4open.science/r/XAI-fool-388D.

1 Introduction

The impact of transformer models in natural language processing has been paramount, and, as such, research on interpretability has also been steadily advancing (Fantozzi and Naldi, 2024). One family of model-specific explanations focuses on feature-level, i.e., token-level, explainability methods that yield saliency or relevance maps (Schneider, 2024). Such explanations are specifically helpful for smaller, task-specific (e.g., classification) models, as they demonstrate granular insight into a model's workings per sample (Liu et al., 2021). However, prior work has demonstrated the instability of such explanations in the image domain, particularly concerning Convolutional Neural Networks (CNNs). Specifically, Heo et al. (2019) demonstrates that explanations can be shaped by an adversary through model manipulation attacks as opposed to, e.g., input-focused attacks via so-called adversarial examples. As such, explainability methods are fooled by training that shapes model explanations through supervised explanation learning.

Explorations of transformer- and NLP-specific cases are scarce, making it unclear whether these vulnerabilities extend to the domain and architecture (Baniecki and Biecek, 2024). We extend the Heo et al.'s (2019) framework to encoder models in the NLP context by proposing new attack scenarios and optimisation approaches. We ask the following questions: Can encoder-based NLP classifiers be trained to manipulate token-level explanations produced by state-of-the-art transformer explainability methods without degrading task performance?

Related Work. Several transformer-specific explainability methods have been proposed. Earlier methods tend to focus on the attention mechanism as an approach to explainability (Fantozzi and Naldi, 2024). More advanced methods incorporate these transformer-specific elements but go beyond them. For instance, Ali et al. (2022) identified the failure modes of GradientxInput approaches (Denil et al., 2014) using the Layer-wise relevance propagation (LRP) framework (Bach et al., 2015) and demonstrated strong conservation properties through their adaptation. In contrast, Chefer et al. (2021a) generalises the Chefer et al.'s (2021b) method by explicitly building on attention-focused methods and extending them via gradient-based weighting, making the explanations class-specific and yielding the Generic Attention Explainability (GAE). Due to their strong empirical performance and widespread use, these methods form the primary targets of our analysis.

To our knowledge, this is the first systematic study of adversarial training-time manipulation of token-level explanations in transformer-based NLP

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 44490–44509
October 24-29, 2026 ©2026 Association for Computational Linguistics

classifiers. Earlier contributions focus on improving faithfulness through explanation-aware training and/or employ auxiliary models (Moradi et al., 2021, 2020; Ismail et al., 2021; Wang et al., 2025; Barkan et al., 2024; Wang et al., 2020; Atanasova et al., 2022), while our contributions explicitly focus on adversarial objectives without additional models that need to be trained. We introduce novel sequence-aware attack objectives that adapt model manipulation approaches to text inputs, including ranking-based formulations, which enable stable manipulation of token relevance. Simultaneously, task accuracy is largely preserved with absolute drops of usually less than 5% and typically 1%. We show that strong transformer-specific explainability methods can exhibit systematic changes in attribution patterns under such training, particularly when suppressing highly influential tokens: previously highly relevant tokens rank among the top 10% least relevant tokens after the attack. Importantly, we highlight that these changes are in some cases accompanied by shifts in the model's reliance on alternative input signals, rather than merely misleading a specific explanation method, as suggested by Heo et al. (2019). Finally, we find that successful manipulations can generalise across explainability methods, indicating shared structural weaknesses. Our findings provide insight into the interaction between explanation methods and learned decision strategies, highlighting the need to consider model adaptation when interpreting explanation robustness.

2 Methodology

We study a training-time adversary who controls model fine-tuning but not inference. Given training data D, initial parameters w_0, and final parameters w, the adversary modifies the fine-tuning objective to shape token-level explanations while preserving task accuracy. The adversary does not alter test inputs, predictions, or explanations at inference time and does not assume knowledge of future test samples or of the explanation method used during deployment, which aligns with related work (Heo et al., 2019).

2.1 Dual Loss Framework and Optimisations

Model manipulation attacks are implemented using a dual-loss approach (Heo et al., 2019):

$$\mathcal{L}(D, D_{\text{fool}}, I; w, w_0) = \mathcal{L}_C(D; w) + \lambda \mathcal{L}_F^I(D_{\text{fool}}; w, w_0), \quad (1)$$

where L_C is cross-entropy and L_F^I penalises explanations produced by method I. The subset $D_{\text{fool}} \subseteq D$ contains samples affected by the attack, and λ controls the trade-off between task performance and explanation manipulation.

We instantiate L_F^I with value-based, distributional, and ranking-based objectives. Since token-level explanations are commonly interpreted by relative importance, our main analysis emphasises ranking-based objectives. Full loss definitions are provided in Appendix D.

We consider two manipulation goals. **Elevation** increases the relevance of uninformative target tokens, selected either by position or token identity. This tests whether arbitrary or weakly informative tokens can be made salient without changing the prediction task. **Suppression** decreases the relevance of tokens that are highly ranked by a benign reference model, either per sample or across the training corpus. This tests whether influential tokens can be demoted while preserving classification performance. These goals define the construction of D_{fool} and the target masks used by L_F^I.

2.2 Experimental Evaluation

We evaluate whether explanation-aware training changes token saliency while preserving accuracy, and whether such changes correspond to altered model behaviour. We compare attacked and benign models in terms of task accuracy, attacked-token rank, cross-method transfer, and masked-token performance to explore functional reliance on attacked tokens.

For elevation, we target either the first non-special token or frequent tokens that are not highly relevant under the benign model. For suppression, we target either the top-ranked token per sample or the tokens most frequently ranked as most relevant in the training set. The first setting creates sample-specific attacks, while the second creates corpus-level attacks against recurring salient tokens. D_{fool} consists of the training samples containing the relevant target tokens; token-selection details are given in Appendix F.

All experiments are performed on four datasets and four models using first-order only optimisations. Simple binary sentiment classification tasks are chosen next to more complex natural language inference tasks on the model's attribution behaviour. Accordingly, the datasets are the SST-2 and IMDb (Maas et al., 2011; Socher et al., 2013) as well as FEVER and e-SNLI (Thorne et al., 2018; Camburu et al., 2018), for which statistics are presented in Appendix A. We test all approaches with BERT (Devlin et al., 2019), ModernBERT (Warner et al., 2024), DistilBERT (Sanh et al., 2020), and ALBERT (Lan et al., 2020) to explore the effects of different architectures via weight sharing. All experiments are implemented using PyTorch (Paszke et al., 2019) and Transformers (Wolf et al., 2020) and run on NVIDIA GeForce RTX 2080 Ti GPUs. Hyperparameters and further details are available in Appendix F. Given that second-order gradients are unstable and memory-intensive (Kashyap, 2023), we treat the gradients introduced by the explanation pass as constants in the optimisation procedure. Taken together, this combination allows us to examine robustness across datasets and architectures without overspecifying the attack setting.

The primary evaluation metric is the average ranking percentile of attacked tokens, reflecting the common interpretation of token relevance as a ranking (Liu et al., 2021). For each occurrence of a token of interest, we compute its percentile rank within the relevance distribution of the corresponding sequence and average these percentiles across all occurrences in D_{fool}. Higher values indicate successful elevation, whereas lower values indicate successful suppression.[1]

Evaluation is performed based on the validation set as test sets are typically held out. Standard hyperparameters, such as the learning rate, were not tuned (see Appendix F). However, the attack strength is tuned and later fixed for additional models and datasets. Concretely, AlBERT and BERT sentiment classification tasks served as the primary attack strength tuning cases, leading to potential overfitting of the attack to those datasets.

Finally, we analyse how explanation-aware training affects functional model behaviour beyond attribution metrics. We mask attacked tokens in the validation set to assess whether attribution changes coincide with shifts in model reliance, and evaluate whether these changes persist when explanations are generated using an alternative method not used during training.

3 Results & Discussion

Table 1 summarises the main findings for BERT on SST-2, with several consistent trends observed across additional settings reported in Appendix B. All reported results use $\lambda = 10$, which we found to provide a stable balance between task accuracy and attribution manipulation.[2] We evaluate effects on (i) core task accuracy, (ii) changes in token relevance rankings, (iii) transfer across explainability methods, and (iv) functional model behaviour.

Accuracy preservation. Across all settings, explanation-aware training has a limited impact on task performance, with absolute accuracy drops of at most 1–5%, with few exceptions on the FEVER dataset (see Appendix B). This indicates that models can accommodate explanation-aware constraints without substantially degrading predictive performance. Thus, attribution changes are model adaptations rather than artefacts of task failure. Nevertheless, label flipping can occur, as the example in Appendix C shows.

Token Saliency Changes. Manipulation effectiveness differs markedly by goal and model. Relevance suppression is more effective than relevance elevation, with highly ranked tokens consistently demoted to lower percentiles, especially for the LRP-trained models.[3] In contrast, elevation only achieves comparable performance in some settings when trained on GAE.[4] We observe further differences among the explainability methods, especially regarding cross-method generalisation. Only the LRP-trained manipulations based on BERT generalise well for token suppression, whereas other settings and models tend to generalise inconsistently. These asymmetries suggest varying model adaptations across explainability methods and models (see Appendix B for details).

These findings are also supported by qualitative explorations. For instance, Figure 1 shows how frequently relevant tokens can be decreased in relevance. Notably, “a" and “," are also highlighted as they appear as overall highly relevant tokens due to their unusually frequent appearance. Sim-

[1] For details, see Appendix F.

[2] For a sensitivity analysis, see Appendix D.

[3] Only a few exceptions show minimal margins.

[4] Notably, GAE appears more susceptible to elevation attacks, as cross-evaluation of across multiple LRP-attacked models shows, while the LRP-based evaluation shows moderate success at best. See Appendix B for details.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	91.71	59.43	67.72	57.70	0.83 / 0.81
	Rank Loss	**92.52**	**66.76**	**68.41**	**59.60**	0.82 / 0.82
	Baseline	50.96	52.17	52.17	50.96	0.83 / 0.83
Elevating specific tokens	MSE	**83.06**	53.58	50.29	26.66	0.82 / 0.81
	Rank Loss	80.45	**59.74**	**55.16**	**31.84**	0.82 / 0.81
	Baseline	41.26	49.01	48.20	30.35	0.82 / 0.82
Decreasing top token per sample	Top-k Loss	40.55	69.31	91.53	83.58	0.82 / 0.80
	Rank Loss	**34.71**	**52.32**	**90.24**	**81.28**	0.82 / 0.81
	Baseline	91.51	82.67	94.77	87.10	0.83 / 0.83
Decreasing overall top token	Top-k Loss	**9.53**	52.83	80.55	65.18	0.82 / 0.82
	Rank Loss	14.00	**18.04**	**77.19**	**57.15**	0.82 / 0.82
	Baseline	85.20	58.97	89.12	74.15	0.83 / 0.83

Table 1: Mean percentile rank comparison across explainability methods for the BERT model on SST-2. Columns report mean percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and mean task accuracy (GAE / LRP). Results are averaged over three training runs and evaluated on the validation set. Standard deviations are omitted. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.826	0.839/0.820	0.748/0.780
LRP	0.826	0.838	0.898
GAE	0.817	0.857	0.777

Table 2: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE).

ilar effectiveness can be found for location-based attacks, which further underlines the quantitative results (see Appendix C).

Ranking-based objectives typically outperform magnitude- and distribution-based alternatives. Unlike value-based losses, ranking constraints directly operate on relative token importance, which aligns with how feature-level explanations are typically interpreted and compared (see, for instance, Ali et al. (2022); Liu et al. (2021)). Additional comparisons are provided in Appendix E.

Model Behaviour. Masking experiments are consistent with the hypothesis that explanation-aware training can alter model reliance rather than merely misleading the explainability method. Under LRP-based training, masking attacked tokens improves and otherwise decreases accuracy, as shown in Table 2. Further results in Appendix B partially support these findings with significant differences between models. Generally, models tend to be more robust under successful token suppression attacks, indicating less reliance on targeted tokens.

This finding suggests that the model learns to treat attacked tokens as noise. In contrast, accuracy decreases under GAE-based training, indicating that the model continues to rely on these tokens despite reduced attribution, albeit often to a lesser degree, as the results in Appendix B show. Our optimisation procedure offers a plausible explanation for the observed differences in model adaptation and cross-method generalisation. In our setup, optimisation is restricted to first-order gradients. In LRP-based training, these gradients propagate only to the input embeddings, whereas in GAE-based training, they also influence the attention mechanism.

We hypothesise that when suppressing token relevance, LRP-based training therefore appears to introduce noise at the embedding level, leading the model to downweight previously dominant tokens without directly modifying the attribution mechanism. This behaviour is consistent with the masking results: performance can improve when such tokens are masked, suggesting that the model had over-relied on noisy representations and can fall back on alternative decision cues. In contrast, GAE-based training more directly constrains attribution scores via attention, resulting in reduced apparent relevance while often preserving reliance on the same tokens, and consequently lower masked accuracy. Future research can investigate the specific mechanistic influence of the training procedure to understand the impact of token-level reliance in further detail and confirm the hypothesised mechanism, as well as understand the partial confirma-

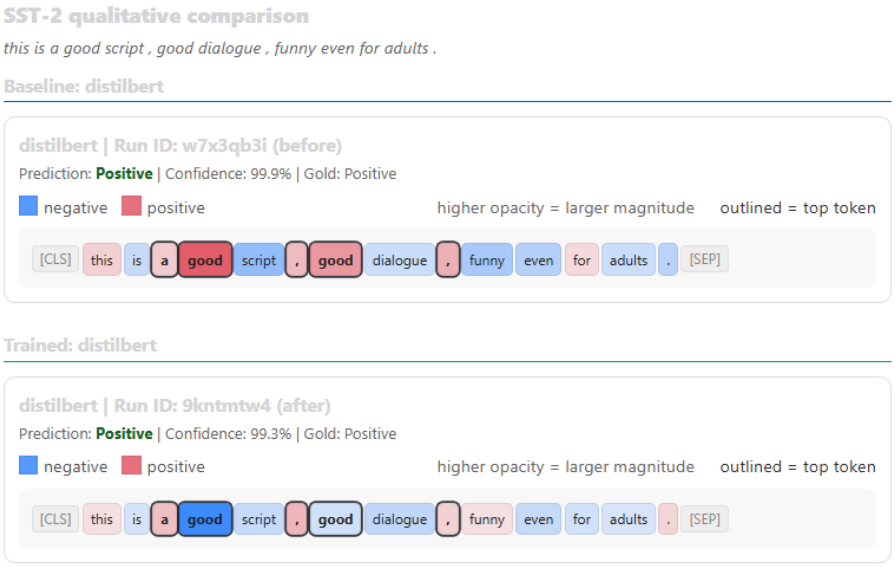


Figure 1: Qualitative example for a DistilBERT model attacked with LRP on SST-2. The aim of the attack is to decrease the relevance of tokens that appeared often as the most important ones. These tokens are highlighted.

tions detailed in Appendix B.

This hypothesis helps explain cross-method generalisation. Embedding-level adaptations induced by LRP-based training propagate throughout the model and affect attribution behaviour under alternative explainability methods.[5] In contrast, attention-focused adaptations are more tightly coupled to the specific explanation mechanism used during training. This perspective clarifies why relevance elevation is particularly difficult under LRP-based training: artificially amplifying token importance would require either increasing embedding dominance or suppressing competing tokens, conflicting with maintaining task performance.

Taken together, our results show that explanation-aware training in encoder-based NLP models does not uniformly "fool" explainability methods, but could instead induce systematic changes in how models distribute reliance across input features. We find that these changes depend on the interaction between the explanation method and the optimisation objective, with embedding-level adaptations exhibiting stronger cross-method generalisation than attention-focused constraints. Across settings, ranking-based objectives emerge as a particularly effective and stable mechanism for manipulating token-level explanations for adversarial purposes, as they directly constrain relative token importance in a manner aligned with how feature-level explanations are interpreted in language models. These findings extend Heo et al.'s (2019) prior work on model manipulation to the transformer domain and highlight that apparent explanation vulnerabilities often reflect deeper model adaptations.

Practical implications. These findings suggest that token-level explanations should not be audited only as post-hoc visualisations of fixed models. Under explanation-aware training, changing the explanation objective can also change the learned decision strategy. Auditing pipelines should therefore combine attribution checks with functional-behavioural tests, such as cross-method evaluation and token masking, especially when explanations are used to assess model reliance.

Limitations and Ethical Concerns

Our study focuses on a narrow experimental setting: binary sentiment classification and two natural language inference tasks using transformer-based classifiers, English-language datasets, and a limited set of token-level explainability methods. While this choice allows for controlled analysis of explanation-aware training, it restricts the generalisability of our findings to other languages, tasks (e.g., sequence labelling or generation), model families, or explanation paradigms. Moreover, we restrict optimisation to first-order gradients and eval-

[5]This generalisation does not manifest in ALBERT models and could be related to the weight sharing and thus repetitive strong influence of the same weights on the explanation.

uate a small number of datasets and architectures without conducting an extensive hyperparameter search, prioritising interpretability and analytical clarity over exhaustive optimisation. Due to computational constraints, all reported experiments are single-run evaluations except for Table 1; we therefore interpret consistency across datasets, models, and attack settings as suggestive rather than as a substitute for seed-level uncertainty estimates. Future work is needed to assess whether the observed adaptation dynamics extend across random seeds, to multilingual settings, more complex tasks, larger models, or alternative optimisation regimes.

The techniques studied in this work have a high potential for misuse. Explanation-aware training could be employed to deliberately obscure model reliance on sensitive or undesirable input features, thereby undermining the use of explainability methods for auditing, accountability, or regulatory compliance. While our goal is to diagnose and understand limitations of token-level explanations, we emphasise that such methods should not be used to evade oversight. We therefore present our results as a cautionary analysis, highlighting the need for robustness-aware interpretability and complementary evaluation strategies when explanations are used in high-stakes settings.

Acknowledgments

The authors used generative AI tools in accordance with the ACL Policy on AI Writing Assistance. Specifically, such tools were used for literature search support; for generating and refining low-novelty text describing pre-existing ideas (which were subsequently verified for accuracy and accompanied by appropriate citations); and for exploratory brainstorming, followed by independent literature review and validation by the authors. In addition, AI-based coding assistants (e.g., predictive code completion and conversational debugging support) were used during software development. All scientific decisions, analyses, interpretations, and final text remain the sole responsibility of the authors.

References

Ameen Ali, Thomas Schnake, Oliver Eberle, Grégoire Montavon, Klaus-Robert Müller, and Lior Wolf. 2022. XAI for transformers: Better explanations through conservative propagation. In *International conference on machine learning*, pages 435–451. PMLR.

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2022. Diagnostics-guided explanation generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10445–10453.

Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140.

Hubert Baniecki and Przemyslaw Biecek. 2024. Adversarial attacks and defenses in explainable artificial intelligence: A survey. *Information Fusion*, 107:102303.

Oren Barkan, Yonatan Toib, Yehonatan Elisha, Jonathan Weill, and Noam Koenigstein. 2024. LLM explainability via attributive masking learning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9522–9537, Miami, Florida, USA. Association for Computational Linguistics.

Oana-Maria Camburu, Tim Rockt"aschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9539–9549. Curran Associates, Inc.

Hila Chefer, Shir Gur, and Lior Wolf. 2021a. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 397–406.

Hila Chefer, Shir Gur, and Lior Wolf. 2021b. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 782–791.

Misha Denil, Alban Demiraj, and Nando De Freitas. 2014. Extraction of salient sentences from labelled documents. *arXiv preprint arXiv:1412.6815*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Paolo Fantozzi and Maurizio Naldi. 2024. The explainability of transformers: Current status and directions. *Computers*, 13(4):92.

Juyeon Heo, Sunghwan Joo, and Taesup Moon. 2019. Fooling neural network interpretations via adversarial model manipulation. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Aya Abdelsalam Ismail, Hector Corrada Bravo, and Soheil Feizi. 2021. Improving deep learning interpretability by saliency guided training. *Advances in Neural Information Processing Systems*, 34:26726–26739.

Rohan Kashyap. 2023. A survey of deep learning optimizers – first and second order methods. *Preprint*, arXiv:2211.15596.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *International Conference on Learning Representations*.

Shengzhong Liu, Franck Le, Supriyo Chakraborty, and Tarek Abdelzaher. 2021. On exploring attention-based explanation for transformer models in text classification. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 1193–1203.

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 *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.

Pooya Moradi, Nishant Kambhatla, and Anoop Sarkar. 2020. Training with adversaries to improve faithfulness of attention in neural machine translation. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 93–100.

Pooya Moradi, Nishant Kambhatla, and Anoop Sarkar. 2021. Measuring and improving faithfulness of attention in neural machine translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2791–2802, Online. Association for Computational Linguistics.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, and 2 others. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *Preprint*, arXiv:1910.01108.

Johannes Schneider. 2024. Explainable generative AI (GenXAI): a survey, conceptualization, and research agenda. *Artificial Intelligence Review*, 57(11):289.

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.

James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *NAACL-HLT*.

Junlin Wang, Jens Tuyls, Eric Wallace, and Sameer Singh. 2020. Gradient-based analysis of NLP models is manipulable. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 247–258. Association for Computational Linguistics.

Yifan Wang, Sukrut Rao, Ji-Ung Lee, Mayank Jobanputra, and Vera Demberg. 2025. B-cos lm: Efficiently transforming pre-trained language models for improved explainability. *arXiv preprint arXiv:2502.12992*.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *Preprint*, arXiv:2412.13663.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. HuggingFace's Transformers: State-of-the-art natural language processing. *Preprint*, arXiv:1910.03771.

A Dataset Statistics

Table 3 shows the statistics for the binary sentiment classification task. Table 4 shows the statistics for the multi-class Natural-Language Inference (NLI) tasks.

B Additional Results

B.1 Attack Effectiveness

This appendix presents results for the additional dataset and model combinations. We do not show

Statistic	SST-2	IMDb
Training size	77,300	25,000
Validation size	872	12,500
Sequence Length:		
Median	10 (24)	233 (232)
Mean	13.3 (25.2)	313.9 (309.5)
Min.	3 (4)	13 (10)
Max.	66 (55)	3127 (3157)
Positives Rate	0.56	0.5

Table 3: Dataset statistics for SST-2 and IMDb. The IMDb validation set is created by randomly sampling 50% of the original test set. Length statistics are based on the bert-base-uncased tokeniser for the training set, with validation statistics in parentheses (pre-truncation).

Statistic	FEVER	e-SNLI
Training size	145,449	549,367
Validation size	19,998	9,842
Sequence Length:		
Median	17 (18)	25 (26)
Mean	20.2 (20.7)	26.2 (27.5)
Min.	8 (8)	7 (9)
Max.	362 (192)	125 (113)
Class Distribution (train):		
Class 0	0.20	0.33
Class 1	0.55	0.33
Class 2	0.25	0.33

Table 4: Dataset statistics for FEVER and e-SNLI. Length statistics are based on the bert-base-uncased tokeniser for the training set, with validation statistics in parentheses (pre-truncation). Validation class distribution is balanced. For FEVER: Class 0 = Refuted, Class 1 = Supports, Class 2 = NotEnoughInfo (sampled second part from documents, as in Thorne et al. (2018)). For e-SNLI: class 0 = entailment, class 1 = neutral, class 2 = contradiction. Samples are constructed by concatenating the premise, the hypothesis/claim, and the evidence.

all combinations exhaustively here; instead, we display representative results such that we show at least one classification task and one NLI task per model.

The additional results also include experiments on functional reliance. Here, we focus on the suppression setting of overall relevant tokens, as this setting provides the clearest test of whether the model learns to reduce reliance on tokens that were previously stable predictors.

Table 5 highlights the continued strong performance of LRP-based token suppression for the BERT model. However, the performance of GAE-based training is weaker than on SST-2, indicating that more extended sequences pose challenges for clear optimisation. Additionally, Table 6 shows that the general trends also hold for more complex tasks like natural language inference. However, the task accuracy is hardly above the majority class share, making the results less indicative of the dual learning goal performance.

Table 7 shows the results for the ALBERT model on a sentiment classification task. Again, the table underlines the strong performance of LRP-based token suppression, albeit with poorer evidence of generalisation, likely due to the attention mechanism's consistent impact through weight-sharing. They also highlight further difficulties of attacks based on GAE, potentially for similar reasons. While the attacks do elevate or suppress target tokens, the magnitude of change, i.e., attack effectiveness, is lower than for BERT-based experiments. However, generalisation appears stronger, indicating LRP's high sensitivity to attention weights.

Table 8 shows the results for AlBERT on the FEVER dataset. Overall, the results show successful token elevation under GAE attacks and successful token suppression, especially for LRP, in line with the above-shown results. Moreover, we find some evidence for generalisation and successful token elevation under LRP, further demonstrating attack effectiveness on more complex datasets while maintaining stronger task performance than the BERT counterpart. Yet, it is noteworthy that accuracy tends to collapse under token suppression using LRP and top-k loss.

Tables 9 and 10 show the attack effectiveness metrics for ModernBert on SST-2 and e-SNLI. Some general trends such as strong elevation under GAE-based attacks as well as weaker effectiveness under LRP-attacks are confirmed. Overall, token suppression is less effective for this model, yet some suppression can still be identified. Highly relevant tokens across the dataset are typically ranking around the 50th percentile and are not strongly decreased in their relevance, which is to be expected given their already moderate aggregate relevance. Also, generalisation is moderately effective for LRP-based elevation attacks that are evaluated using GAE. At the same time, generalisation is very weak of absent for previously identified effective settings like token suppression.

Finally, Tables 11 and 12 further underline that the principal findings largely hold for additional models, also in more complex datasets such as e-SNLI, another natural language inference dataset. However, we also see that token suppression for overall tokens is less effective, possibly due to inconsistent signals, as their baseline relevance is rather low. Contrastingly, LRP-based elevation

tends to work well in some of the settings for DistilBERT, including some successful generalisation.

Overall, we see a mixed picture. Generally, GAE-based elevation tends to work well, and also token suppression, particularly LRP-based suppression if the tokens are consistently high-ranked. Cross-generalisation of attacks is rather weak and can occur model-dependent with LRP-based elevation, especially for location-based attacks and in some cases also for suppression.

B.2 Token Reliance

Token reliance experiments indicate a mixed picture across additional models and dataset. The AlBERT model on IMDb shows the clearest confirmation of the above results: accuracy under token masking improves relative to the baselines for LRP-attacked models while the accuracy drops under GAE, as can be seen in Table 17.

For other model settings, the findings can partially be confirmed. For instance, Table 13 shows that accuracy recovers slightly under token masking for the LRP-attacked model relative to the unattacked model. However, performance still remains below the majority class share and below that of the same model without token masking, indicating that tokens are not treated as noise but that reliance is diminished. Similar results can be seen in Tables 15 and 19, where the attacked models are less vulnerable to token masking, especially the LRP-attacked model.

Likewise, Tables 16 and 14 show that both attacked models are less vulnerable to token masking. However, the GAE-attacked model less so, contradicting the above-presented hypothesis. Finally, ModernBERT on e-SNLI indicates complete collapse of performance under token masking.[6]

Overall, we tend to see less vulnerability of attacked models, particularly for LRP-based models, although the effect relative to the GAE attacked model is often minimal. These findings, however, highly depend on the specific model attacked and the effectiveness of the attack itself.

C Additional Qualitative Example

This section displays an additional qualitative example that underlines the overall quantitative findings presented above. We focus on examples from the SST-2 dataset as these are highly interpretable and short, making them suitable for display.

Figure 2 shows the successful attack of elevating the relevance of the token in first position, while a baseline model and common human interpretation would not attribute much importance to the token in first position. However, in this case we also witness the change of the model classification, indicating more profound training procedure impacts. This example also highlights that accuracy preservation does not necessarily indicate sample-level stability.

D Loss Function Formulations

Generally, the adversary's goal can be pursued by optimising the objective function below, adopted from (Heo et al., 2019):

$$L(D, D_{\text{fool}}, I; w, w_0) = L_C(D; w) + \lambda \mathcal{L}_F^I(D_{\text{fool}}; w, w_0), \quad (2)$$

where

- $\mathcal{L}_C$ denotes the standard loss function, in this case, cross-entropy loss.
- $\mathcal{L}_F^I$ denotes the penalty imposed based on the explanation, specific to
- D_{fool}, the dataset used to optimise the attack, and potentially sub-sampled from
- D, the original training set used to finetune, some pre-trained model with weights
- w_0 to reach a model with weights w.
- λ controls the attack strength by regulating the balance between the standard loss $\mathcal{L}_C$ and the attack loss $\mathcal{L}_F^I$.

$\mathcal{L}_F^I$ can be adapted to meet a number of requirements. In the following, we detail multiple approaches to increase the relevance of specific tokens, decrease the relevance of specific tokens, or generally learn specific, uninformative distributions as explanations.

D.1 Increasing Token Relevance

Token relevance can be increased by targeting either specific tokens or specific positions of the input. Both approaches can be optimised via a number of different losses, which are shown below.

D.1.1 Mean Squared Error

Mean Squared Error (MSE) is shown in Eq. (3) and is a common loss function for regression tasks. However, it can also be used to reduce the distance between an explanation h_c^I and a target mask m for a given sample of length T, as shown in (Heo

[6] Note that the results of Tables 14 and 18 are based on a hardly effective attack.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	41.44	49.31	48.07	40.76	0.87 / 0.87
	Rank Loss	**99.14**	**53.30**	**56.97**	**45.61**	0.87 / 0.87
	Baseline	40.65	48.58	48.58	40.65	0.88
Elevating specific tokens	MSE	47.36	52.15	51.63	44.97	0.87 / 0.87
	Rank Loss	**71.32**	**52.66**	**59.68**	**45.99**	0.87 / 0.87
	Baseline	42.03	48.79	52.26	42.51	0.88
Decreasing top token per sample	Top-k Loss	**76.89**	66.78	**93.60**	84.46	0.87 / 0.87
	Rank Loss	81.21	**42.97**	93.67	**79.70**	0.87 / 0.87
	Baseline	100.00	100.00	100.00	100.00	0.88
Decreasing overall top tokens	Top-k Loss	**32.98**	48.62	74.96	55.05	0.87 / 0.87
	Rank Loss	36.02	**32.03**	**74.86**	**49.47**	0.87 / 0.87
	Baseline	91.58	88.35	94.45	84.04	0.88

Table 5: Percentile rank comparison across explainability methods for the BERT model on IMDb. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	68.32	50.47	**52.81**	39.14	0.61 / 0.61
	Rank Loss	**94.16**	**62.68**	50.74	**54.61**	0.57 / 0.61
	Baseline	39.50	49.50	49.50	39.50	0.62
Elevating specific tokens	MSE	85.07	61.51	**51.15**	39.24	0.60 / 0.61
	Rank Loss	**87.88**	**70.18**	48.32	**53.61**	0.59 / 0.61
	Baseline	52.12	51.24	48.07	39.05	0.62
Decreasing top token per sample	Top-k Loss	47.28	64.30	**81.99**	79.15	0.60 / 0.59
	Rank Loss	**47.13**	**34.48**	82.90	**70.84**	0.60 / 0.59
	Baseline	100.00	100.00	94.35	100.00	0.62
Decreasing overall top token	Top-k Loss	**13.30**	48.26	**53.23**	48.14	0.61 / 0.61
	Rank Loss	13.68	**34.09**	54.09	**36.15**	0.61 / 0.60
	Baseline	52.71	64.71	63.02	56.97	0.62

Table 6: Percentile rank comparison across explainability methods for the BERT model on FEVER. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	76.08	52.87	52.66	78.58	0.90 / 0.90
	Rank Loss	**99.17**	**64.07**	**54.78**	**96.60**	0.90 / 0.89
	Baseline	73.71	52.47	52.47	73.71	0.90
Elevating specific tokens	MSE	57.86	**52.87**	**54.84**	**38.82**	0.90 / 0.90
	Rank Loss	**58.84**	51.28	49.63	35.29	0.91 / 0.89
	Baseline	50.67	48.36	51.81	37.25	0.90
Decreasing top token per sample	Top-k Loss	94.99	47.86	**25.11**	92.95	0.90 / 0.90
	Rank Loss	**82.63**	**13.00**	53.10	**88.66**	0.90 / 0.89
	Baseline	100.00	100.00	100.00	100.00	0.90
Decreasing overall top tokens	Top-k Loss	59.11	53.37	62.26	**56.59**	0.90 / 0.90
	Rank Loss	**53.63**	**11.72**	**61.78**	86.13	0.90 / 0.88
	Baseline	62.66	68.35	54.74	64.34	0.90

Table 7: Percentile rank comparison across explainability methods for the ALBERT model on IMDb. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

et al., 2019). Unlike images, samples in NLP tasks vary in length. As such, the aggregation method

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	85.62	62.45	**57.57**	68.30	0.71 / 0.68
	Rank Loss	**90.78**	**70.09**	56.38	**75.59**	0.71 / 0.69
	Baseline	44.57	52.61	52.61	44.57	0.71 / 0.71
Elevating specific tokens	MSE	56.65	55.93	51.50	23.86	0.71 / 0.69
	Rank Loss	**72.19**	**64.55**	**52.85**	**37.93**	0.71 / 0.69
	Baseline	35.23	53.57	52.09	36.31	0.71 / 0.71
Decreasing top token per sample	Top-k Loss	41.10	52.66	**50.29**	**57.43**	0.71 / 0.54
	Rank Loss	**36.76**	**13.39**	55.17	65.87	0.71 / 0.70
	Baseline	100.00	100.00	100.00	100.00	0.71 / 0.71
Decreasing overall top token	Top-k Loss	46.82	54.83	49.69	**47.05**	0.71 / 0.65
	Rank Loss	**37.15**	**15.17**	**43.91**	64.71	0.71 / 0.70
	Baseline	70.56	59.52	48.42	73.39	0.71 / 0.71

Table 8: Percentile rank comparison across explainability methods for the AlBERT model on FEVER. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	91.71	54.48	49.78	76.64	0.78 / 0.79
	Rank Loss	**94.15**	**55.66**	50.85	**86.41**	0.79 / 0.77
	Baseline	56.43	51.04	**51.04**	56.43	0.79
Elevating specific tokens	MSE	83.23	53.26	**53.86**	44.18	0.79 / 0.78
	Rank Loss	**85.94**	**55.77**	52.48	**51.09**	0.79 / 0.79
	Baseline	59.54	51.43	53.66	49.66	0.79
Decreasing top token per sample	Top-k Loss	93.75	58.68	53.80	**91.00**	0.78 / 0.77
	Rank Loss	**62.61**	**55.12**	**53.25**	92.74	0.78 / 0.77
	Baseline	100.00	100.00	100.00	100.00	0.79
Decreasing overall top token	Top-k Loss	55.01	54.71	**51.43**	**62.62**	0.80 / 0.78
	Rank Loss	**46.16**	**50.34**	52.45	72.80	0.79 / 0.80
	Baseline	67.39	52.80	52.60	68.41	0.79

Table 9: Percentile rank comparison across explainability methods for the ModernBERT model on SST-2. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

across N samples can be crucial. Hence, we propose experimenting with both micro- and macro-averaging.

$$\mathrm{MSE}_i = \frac{1}{T_i} \sum_{t=1}^{T_i} \left(h^I_{c_i,t}(w, x_i) - m_{i,t} \right)^2 . \quad (3)$$

MSE Loss - Macro

Macro-averaging as shown in Eq. (4) implies that each sample bears equal weight in the size of the penalty and has the advantage that longer samples do not exert overly strong influence over the overall learning process.

$$\begin{aligned} \mathcal{L}^I_{\mathrm{MSE}_{\mathrm{macro}}} &= \frac{1}{N} \sum_{i=1}^{N} \mathrm{MSE}_i \\ &= \frac{1}{N} \\ &\quad \sum_{i=1}^{N} \left(\frac{1}{T_i} \sum_{t=1}^{T_i} \left(h^I_{c_i,t}(w, x_i) - m_{i,t} \right)^2 \right) . \end{aligned} \quad (4)$$

MSE Loss - Micro

Contrastingly, micro-averaging as shown in Eq. (5) allows for fine-grained feedback per token, allowing the relevance of tokens in long sequences to not be suppressed.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	92.58	66.78	50.98	85.05	0.80 / 0.77
	Rank Loss	**95.44**	**75.22**	51.61	**83.59**	0.80 / 0.79
	Baseline	78.18	55.40	**55.40**	78.18	0.80
Elevating specific tokens	MSE	**81.77**	65.60	50.71	**65.50**	0.80 / 0.77
	Rank Loss	81.36	**69.50**	**51.04**	57.22	0.80 / 0.79
	Baseline	45.37	51.93	49.73	57.26	0.80
Decreasing top token per sample	Top-k Loss	71.79	56.21	56.94	89.10	0.80 / 0.79
	Rank Loss	**47.39**	**39.05**	**55.01**	**84.10**	0.80 / 0.79
	Baseline	100.00	100.00	100.00	100.00	0.80
Decreasing overall top token	Top-k Loss	49.81	53.10	59.00	62.29	0.80 / 0.79
	Rank Loss	**34.43**	**43.05**	56.70	62.78	0.79 / 0.79
	Baseline	74.54	51.90	**52.34**	**58.09**	0.80

Table 10: Percentile rank comparison across explainability methods for the ModernBERT model on e-SNLI. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	95.01	59.14	50.04	51.60	0.84 / 0.83
	Rank Loss	**95.10**	**67.55**	**50.84**	**53.49**	0.81 / 0.84
	Baseline	47.64	50.55	50.55	47.64	0.84
Elevating specific tokens	MSE	76.41	52.39	**46.53**	43.54	0.84 / 0.82
	Rank Loss	**92.39**	**56.28**	44.61	42.28	0.83 / 0.83
	Baseline	64.48	49.75	45.66	**46.23**	0.84
Decreasing top token per sample	Top-k Loss	40.11	83.70	**69.44**	74.75	0.80 / 0.82
	Rank Loss	**32.43**	**72.44**	71.86	**73.08**	0.84 / 0.80
	Baseline	100.00	100.00	100.00	100.00	0.84
Decreasing overall top token	Top-k Loss	28.93	54.30	**53.32**	59.75	0.84 / 0.84
	Rank Loss	**26.68**	**36.43**	52.56	**57.55**	0.85 / 0.83
	Baseline	69.90	59.35	53.52	63.72	0.84

Table 11: Percentile rank comparison across explainability methods for the DistilBERT model on SST-2. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

$$\mathcal{L}^{I}_{\text{MSE}_{\text{micro}}} = \frac{1}{\sum_{i=1}^{N} T_i} \sum_{i=1}^{N} \sum_{t=1}^{T_i} \left(h^{I}_{c_i,t}(w, x_i) - m_{i,t} \right)^2. \quad (5)$$

D.1.2 Categorical Approaches

Given that only a single or very few tokens are emphasised, loss functions originating from more categorical approaches are also suitable candidates. However, unlike purely categorical losses, setting the remaining tokens to zero will likely yield unstable results due to strong gradient feedback. As such, we propose to adapt cross-entropy loss (CE loss) to a harder and softer version of the KL-divergence. Below, we define these two loss functions for this case.

For a sequence x_i of length T_i with class label c_i, let $h^{I}_{c_i}(w, x_i) \in \mathbb{R}^{T_i}$ denote the relevance scores produced by interpretation method I for model parameters w. First, we define a temperature-scaled softmax distribution of relevance scores over tokens:

$$p_{i,t} = \frac{\exp\left(h^{I}_{c_i,t}(w, x_i)/\tau_p\right)}{\sum_{s=1}^{T_i} \exp\left(h^{I}_{c_i,s}(w, x_i)/\tau_p\right)},$$

$$t = 1, \ldots, T_i.$$

In the hard-target case, we normalise the target mask into a categorical distribution:

$$q_{i,t} = \frac{\text{target}_{i,t}}{\sum_{s=1}^{T_i} \text{target}_{i,s}}, \qquad t = 1, \ldots, T_i.$$

Alternatively, in the soft-target case, the target scores are treated similarly to logits and converted

Attack Category	Loss	GAE	LRP	Train GAE → LRP	Train LRP → GAE	Accuracy (GAE / LRP)
Elevating token by position	MSE	**96.03**	65.11	**55.51**	**45.82**	0.84 / 0.84
	Rank Loss	95.67	**69.57**	50.47	36.87	0.80 / 0.83
	Baseline	31.04	51.75	51.75	31.04	0.84
Elevating specific tokens	MSE	77.93	60.12	52.72	44.24	0.84 / 0.84
	Rank Loss	**90.80**	**60.98**	**55.15**	**48.63**	0.83 / 0.84
	Baseline	38.31	52.27	52.22	39.09	0.84
Decreasing top token per sample	Top-k Loss	56.41	63.11	55.51	72.64	0.84 / 0.84
	Rank Loss	**43.73**	**33.14**	**54.41**	**67.80**	0.84 / 0.82
	Baseline	100.00	100.00	100.00	100.00	0.84
Decreasing overall top token	Top-k Loss	31.91	52.07	53.44	**45.57**	0.84 / 0.84
	Rank Loss	**25.26**	**38.38**	54.73	47.10	0.85 / 0.84
	Baseline	52.68	51.47	**53.36**	50.13	0.84

Table 12: Percentile rank comparison across explainability methods for the DistilBERT model on e-SNLI. Columns report percentile ranks under GAE and LRP, cross-generalisation when training attacks on one method and evaluating with the other, and task accuracy (GAE / LRP). All values are single-run validation set evaluations. The best-performing loss function for each approach is highlighted; higher percentile ranks are better for elevation attacks, whereas lower percentile ranks are better for suppression attacks.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.6178	0.6178/0.6178	0.4696/0.4352
LRP	0.6014	0.6014	0.5186
GAE	0.6054	0.6054	0.4539

Table 13: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the FEVER dataset with BERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.7878	0.7878/0.7878	0.4942/0.6548
LRP	0.7982	0.7982	0.5240
GAE	0.7878	0.7878	0.6869

Table 14: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the SST-2 dataset with ModernBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

into a probability distribution via a soft-max function using a temperature τ_q:

$$q_{i,t} = \frac{\exp\big(\text{target}_{i,t}/\tau_q\big)}{\sum_{s=1}^{T_i} \exp\big(\text{target}_{i,s}/\tau_q\big)}, \qquad t = 1, \ldots, T_i.$$

Finally, the loss is computed as the KL divergence between both distributions, with the difference being whether the target distributions have been softened or not:

$$\begin{aligned} \mathcal{L}^{I}_{\text{KL}} &= \frac{1}{N} \sum_{i=1}^{N} \text{KL}(q_i \,\|\, p_i) \\ &= \frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{T_i} q_{i,t} \, \log \frac{q_{i,t}}{p_{i,t}}. \end{aligned} \tag{6}$$

Generally, a setup of $\tau_p > \tau_q$ leads to a spikier target distribution q, enforcing a strong signal yet

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.8734	0.8734/0.8734	0.78232/0.8097
LRP	0.8666	0.8666	0.8180
GAE	0.8821	0.8822	0.8217

Table 15: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the IMDb dataset with DistilBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.7125	0.7125/0.7125	0.3653/0.3796
LRP	0.7026	0.7026	0.4060
GAE	0.7123	0.7123	0.5174

Table 16: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the fever dataset with AlBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

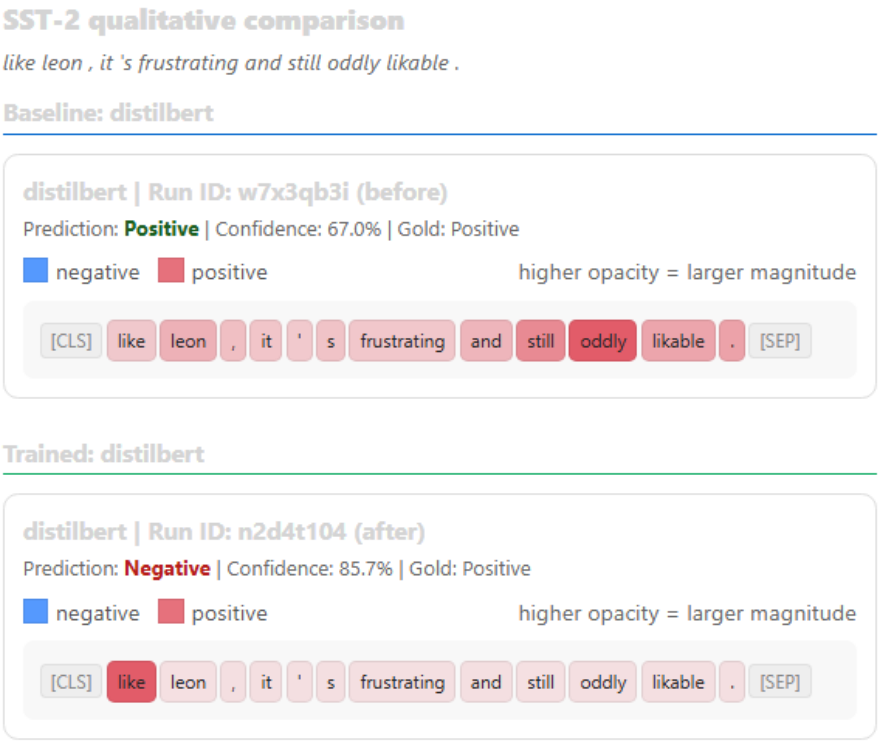


Figure 2: Qualitative example for a model attacked with GAE. The aim of the attack is to emphasise the first token.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.9012	0.9012/0.9012	0.7535/0.8632
LRP	0.8849	0.8848	0.9479
GAE	0.9019	0.9019	0.8465

Table 17: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the IMDb dataset with AlBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.8039	0.8039/0.8039	0.3258/0.3374
LRP	0.7892	0.7892	0.3372
GAE	0.7935	0.7935	0.3392

Table 18: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the e-SNLI dataset with ModernBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

a smoother input distribution p to avoid exploding feedback, especially given that relevance scores will likely not all go to exactly 0, potentially leading the model to react in a stable manner. Accordingly, temperature parameters are set at $\tau_p = 2$ and $\tau_q = 1$, unless otherwise specified. This argument can be extended to the hard case where both the temperature smoothing and softmaxing are entirely absent.

D.1.3 Ranking Approach

Finally, given that only a few tokens are present in some samples, and the main goal is relative ranking rather than overall magnitude, ranking approaches appear suitable. As such, we propose using margin ranking loss as an additional approach to improve token-level relevance.

In this approach, all possible pairs of tokens of a sample i that can be made from the two sets P_i (set of target tokens either to suppress or elevate, depending on y) and N_i (set of remaining tokens) are compared, and a loss is formed using a hinge function as shown in Eq. (7). The parameter m enforces a margin between emphasised and de-emphasised tokens, set at 0.5 for all experiments. Beyond that, the meaning of P_i and N_i can be flipped through y. If $y = 1$, then $p \in P_i$ should be ranked above $n \in N_i$. This setting applies to token elevation experiments. If $y = -1$, then the opposite applies:

Approach	Standard Validation Acc.	Subsampled Validation Acc.	Masked + Subsampled Acc.
Baseline	0.8407	0.8407/0.8407	0.7076/0.7447
LRP	0.8365	0.8365	0.7375
GAE	0.8455	0.8455	0.7682

Table 19: Validation accuracy for standard, subsampled (only samples with overall top tokens), and masked-subsampled evaluation settings for baseline and rank loss approaches on the e-SNLI dataset with DistilBERT. The subsampled dataset only contains samples that contain the top targets (for LRP/GAE). All values are single-run validation set evaluations.

$n \in N_i$ should be ranked above $p \in P_i$, i.e., in the case of token suppression. The resulting loss is normalised by the number of samples N and the number of pairs. Overall, such a loss does not enforce the overall magnitude but rather a token ranking with an additional gap between positive token p and negative tokens n, allowing the model more flexibility yet enforcing the desired explainability characteristics.

$$\mathcal{L}_{\text{rank}} = \frac{1}{\sum_{i=1}^{N} |P_i|\,|N_i|} \sum_{i=1}^{N} \sum_{p \in P_i} \sum_{n \in N_i} \max\Big(0,\ -y\Big(h^I_{c_i,p}(w, x_i) - h^I_{c_i,n}(w, x_i)\Big) + m\Big). \tag{7}$$

All of the approaches above can be "plugged" into the overall objective shown in Eq. (2). Thereby, the tokens whose relevance should be increased are emphasised relative to the relevance of the remaining tokens. By doing so, the approaches cover a wide range of loss functions, such as regression and categorical approaches to ranking loss, allowing for wide exploration. However, they share the commonality that they attack the relevance scores through an overall approach across a given sample by increasing and decreasing a token's desired relevance, depending on the token's status. If successful, such an approach can be used to emphasise any arbitrary token.

D.2 Decreasing Token Relevance

Clearly, the opposite of increasing token relevance is decreasing token relevance. However, in contrast to the above approach, decreasing token relevance should be more targeted, since this attack is explicitly aimed at high-relevance tokens. Hence, the loss functions can accordingly be more targeted than learning entire distributions, as is the case with the categorical and regression approaches from above. Two approaches that aim to fulfil this promise are introduced below.

D.2.1 Top-k Loss

First, simply taking the absolute magnitude of the top k tokens can serve as an adequate penalty to decrease that magnitude. This approach is inspired by the top-k loss of (Heo et al., 2019). However, given the limited number of tokens as compared to pixels in an image, k indicates the number of tokens, rather than the percentage.

$$\mathcal{L}_{\text{Top-}k} = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{k} \sum_{t \in \text{Top-}k\left(h^I_{c_i}(w_0, x_i)\right)} \left| h^I_{c_i,t}(w, x_i) \right|. \tag{8}$$

D.2.2 Ranking Loss

Similar to the approach above, top tokens can be decreased explicitly in terms of ranking. As such, they can be emphasised by setting $y = -1$ and targeting them as part of P_i. Therefore, the optimisation follows the same principles and directly uses Eq. (7).

Both approaches can be applied to different definitions of top tokens. On the one hand, a top token can be understood as the highest-ranking token per sample. On the other hand, it can also be a number of selected tokens, such as the most important tokens across a reference corpus and model. In that case, D_{fool} is subsampled to those samples that contain the relevant tokens.

E Additional Loss Functions & Sensitivity Analysis Results

E.1 Sensitivity Analysis LRP

Tables 20 to 23 show that accuracy tends to decrease as λ increases. This effect becomes noticeable at $\lambda = 10$. Thus, further increases in the attack strength are not advisable to preserve task accuracy. The opposite trend holds for the attack evaluations, with noticeable jumps in losses with higher attack emphasis.

Furthermore, Table MSE macro (see Tables 21 and 20) and top-k (see Tables 22 and 23) losses tend to optimise magnitude-based losses best, whereas ranking loss is typically most efficient for ranking-based objectives.

E.2 Sensitivity Analysis GAE

The results of LRP are mirrored for GAE, as can be seen in Tables 24 and 25 for token elevation, and in Tables 26 and 27 for token suppression.

F Experimental Setup Details

F.1 Hyperparameters

We experiment with a set of standard hyperparameters for comparability. These include a learning rate of 1e-5 with a linear scheduler and 10% warmup steps. The model is optimised for four epochs using AdamW and a batch size of 32, or reduced where necessary, due to memory constraints. These choices are based on previously found good performance for these models and tasks.[7] We use the standard tokenisers, with truncation to a maximum sequence length of 512 from the beginning of each sequence.

Head truncation may underrepresent later evidence. However, evaluation metrics are based on the truncated sequences. The evaluation metrics will thus not be biased upwards for truncated samples. Similarly, token selection is also based on truncated sequences, which aligns token selection and experimental setting. Hence, later tokens are ignored across all settings, potentially leading to imperfect model behaviour, as part of the input is missing, yet not biasing metric evaluations due to the consistent truncation.

All experiments are run with first-order only optimisation, i.e., treating gradients introduced by the explanation pass as constants. In practice, we detach gradients from the explanation's backward pass from the computation graph. All reported experimental results are single runs based on truncated sequences with the given hyperparameters.

ModernBERT and BERT are both truncated due to resource constraints. The procedure for the BERT model is initially implemented by Ali et al. (2022). For ModernBERT, we truncate to 6 blocks.

Under this setup, model training per run takes between 15 minutes (BERT SST-2) and 7 hours for larger models and datasets such as AlBERT on e-SNLI on the described setup. Most setups require 45 minutes to 2.5 hours of compute for a single training run.

[7] See here for a number of different successfully finetuned models https://huggingface.co/textattack/models.

F.2 Details Elevating Token Relevance

Increasing token relevance is split into two experiments: increasing tokens by their position and increasing specific tokens. In terms of position, the token at the first position is chosen, i.e., the first token following the CLS token, resulting in D_{fool} being the full dataset. This token is typically uninformative for sentiment analysis and, therefore, suitable for showing the effectiveness of the attack.

Choosing specific tokens to increase their relevance artificially is primarily guided by the size of D_{fool}. Based on initial tuning, ten tokens are sampled from all tokens that appear in at least 25% of the training samples. Additionally, tokens should not be in the top 10 most relevant tokens per class (see below for details). Suppose only an insufficient number of tokens can be sampled. In that case, the threshold for minimum frequency is lowered, leading to a threshold of 10% or 15% for the SST-2 dataset, depending on the explainability method.[8] The sizes of the resulting attack sets D_{fool} can be found in Table 28. Once a token's position is found, a mask highlighting it can be created. This mask is a binary relevance sequence that can then be used as a target explanation for the optimisation functions shown above.

F.3 Details Decreasing Token Relevance

High-relevance tokens are either identified per sample or across a corpus based on a reference model. In the former case, simply the highest positive token is identified per sample and targeted with the losses detailed above. Thus, D_{fool} is of the same size as the overall dataset. In the latter case, the top token per sample is recorded based on the training set and an unattacked model. The 10 most frequently appearing tokens are saved per class, dataset, model, and tokeniser.[9] The sizes of the resulting attack sets D_{fool} can be found in Table 28. Once the tokens and their positions are identified, a binary mask can be made. In this binary mask, the highly relevant tokens are marked, and their relevance magnitudes are minimised using the optimisation functions above.

[8] Tokens are chosen based on the model-specific tokenisers.

[9] Explanations are with regard to a specific class. We use the target class. As such, the highest-ranking tokens are understood to have the highest positive relevance. Negative relevance implies, in a binary problem, that this token has positive relevance for the opposite class, which is clearly separated here; therefore, the signed magnitude, not the absolute magnitude, is used.

λ	Attack $\mathcal{L}$	Accuracy	MSE Macro	MSE Micro	Rank Loss
0.1	KL Hard	0.830	0.115	0.086	0.508
0.1	KL Soft	0.828	0.115	0.086	0.516
0.1	MSE Macro	0.830	0.112	0.084	0.513
0.1	MSE Micro	0.827	0.115	0.086	0.517
0.1	Rank	0.830	0.115	0.086	0.506
1	KL Hard	0.825	0.130	0.097	0.470
1	KL Soft	0.828	0.110	0.083	0.503
1	MSE Macro	0.820	0.093	0.073	0.492
1	MSE Micro	0.828	0.109	0.082	0.511
1	Rank	0.826	0.115	0.088	0.462
10	KL Hard	0.823	0.204	0.161	0.435
10	KL Soft	0.818	0.092	0.073	0.467
10	MSE Macro	0.805	0.076	0.062	0.465
10	MSE Micro	0.817	0.091	0.070	0.489
10	Rank	0.826	0.096	0.081	0.416

Table 20: *Attack evaluation metrics of the attack on the first token of the sequence by λ and attack loss using LRP.* The table shows the results aiming to increase the relevance of the first token with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	MSE Macro	MSE Micro	Rank Loss
0	-	0.826	0.292	0.272	0.534
0.1	KL Soft	0.827	0.291	0.271	0.533
0.1	MSE Macro	0.829	0.286	0.268	0.531
0.1	MSE Micro	0.828	0.291	0.271	0.533
0.1	Rank	0.826	0.291	0.271	0.530
1	KL Soft	0.829	0.283	0.267	0.530
1	MSE Macro	0.822	0.263	0.255	0.517
1	MSE Micro	0.828	0.282	0.266	0.530
1	Rank	0.817	0.288	0.270	0.504
10	KL Soft	0.818	0.258	0.251	0.508
10	MSE Macro	0.817	0.246	0.244	0.495
10	MSE Micro	0.820	0.262	0.252	0.516
10	Rank	0.819	0.266	0.260	0.462

Table 21: *Attack evaluation metrics of the relevance elevation on specific tokens by λ and attack loss using LRP.* The table shows the results aiming to increase the relevance of the chosen tokens with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	Top k Loss	Rank Loss
0	-	0.815	0.579	0.898
0.1	Rank Loss	0.830	0.519	0.741
0.1	Top k Loss	0.827	0.445	0.765
1	Rank Loss	0.830	0.346	0.478
1	Top k Loss	0.826	0.210	0.629
10	Rank Loss	0.811	0.344	0.395
10	Top k Loss	0.799	0.143	0.580

Table 22: *Attack evaluation metrics of the attack on the highest ranking token of a sequence by λ and attack loss using LRP.* The table shows the results aiming to decrease the relevance of the highest ranking token with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	Top k Loss	Rank Loss
0	-	0.826	0.235	0.608
0.1	Rank Loss	0.823	0.348	0.433
0.1	Top k Loss	0.826	0.171	0.598
1	Rank Loss	0.826	0.504	0.243
1	Top k Loss	0.819	0.098	0.513
10	Rank Loss	0.826	0.621	0.156
10	Top k Loss	0.811	0.071	0.505

Table 23: *Attack evaluation metrics of the attack on the top ten highest ranking tokens of the training set by λ and attack loss using LRP.* The table shows the results aiming to decrease the relevance of these tokens with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	MSE Macro	MSE Micro	Rank Loss
0	-	0.826	0.079	0.064	0.530
0.1	KL Hard	0.827	0.083	0.068	0.533
0.1	KL Soft	0.827	0.083	0.068	0.534
0.1	MSE Macro	0.827	0.083	0.068	0.534
0.1	MSE Micro	0.828	0.083	0.068	0.534
0.1	Rank	0.826	0.083	0.068	0.533
1	KL Hard	0.822	0.082	0.067	0.524
1	KL Soft	0.827	0.083	0.067	0.533
1	MSE Macro	0.826	0.081	0.066	0.530
1	MSE Micro	0.827	0.083	0.067	0.533
1	Rank	0.826	0.081	0.066	0.504
10	KL Hard	0.797	0.075	0.062	0.424
10	KL Soft	0.820	0.079	0.065	0.517
10	MSE Macro	0.833	0.078	0.064	0.501
10	MSE Micro	0.827	0.079	0.065	0.526
10	Rank	0.804	0.077	0.063	0.430

Table 24: *Attack evaluation metrics of the attack on the first token of the sequence by λ and attack loss using GAE.* The table shows the results aiming to increase the relevance of the first token with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	MSE Macro	MSE Micro	Rank Loss
0	-	0.826	0.079	0.064	0.530
0.1	KL Hard	0.825	0.263	0.255	0.444
0.1	KL Soft	0.827	0.263	0.255	0.444
0.1	MSE Macro	0.825	0.263	0.255	0.444
0.1	MSE Micro	0.827	0.263	0.255	0.444
0.1	Rank	0.827	0.263	0.255	0.444
1	KL Hard	0.821	0.262	0.254	0.443
1	KL Soft	0.821	0.263	0.255	0.444
1	MSE Macro	0.817	0.261	0.253	0.440
1	MSE Micro	0.822	0.263	0.254	0.444
1	Rank	0.819	0.260	0.252	0.439
10	KL Hard	0.808	0.253	0.247	0.427
10	KL Soft	0.822	0.258	0.251	0.436
10	MSE Macro	0.818	0.252	0.246	0.425
10	MSE Micro	0.827	0.257	0.250	0.433
10	Rank	0.819	0.251	0.245	0.424

Table 25: *Attack evaluation metrics of the relevance elevation on specific tokens by λ and attack loss using GAE.* The table shows the results aiming to increase the relevance of the chosen tokens with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	Top k Loss	Rank Loss
0	-	0.826	0.046	0.515
0.1	Rank Loss	0.822	0.044	0.513
0.1	Top k Loss	0.826	0.044	0.513
1	Rank Loss	0.815	0.027	0.497
1	Top k Loss	0.813	0.026	0.497
10	Rank Loss	0.819	0.000	0.473
10	Top k Loss	0.820	0.000	0.475

Table 26: *Attack evaluation metrics of the attack on the highest ranking token of a sequence by λ and attack loss using GAE.* The table shows the results aiming to decrease the relevance of the highest ranking token with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the validation set.

λ	Attack $\mathcal{L}$	Accuracy	Top k Loss	Rank Loss
0.1	Rank Loss	0.828	0.019	0.490
0.1	Top k Loss	0.827	0.018	0.490
1	Rank Loss	0.818	0.008	0.480
1	Top k Loss	0.823	0.007	0.480
10	Rank Loss	0.817	0.000	0.474
10	Top k Loss	0.808	0.000	0.475

Table 27: *Attack evaluation metrics of the attack on the top ten highest ranking tokens of the training set by λ and attack loss using GAE.* The table shows the results aiming to decrease the relevance of these tokens with the BERT model on SST-2 with standard hyperparameters as described above. Results are shown for the full validation set.

F.4 Evaluation Metric Details

We compute ranking percentiles at the level of token occurrences. For each sequence, we identify all tokens belonging to a predefined set of tokens of interest. For each such token occurrence (i, j), we compute its percentile rank within the same sequence as

$$p_{i,j} = \frac{100}{|\mathcal{T}_i|} \sum_{k \in \mathcal{T}_i} \mathbb{I}[r_{i,k} \leq r_{i,j}]$$

where $r_{i,k}$ denotes the relevance score of token j in sequence i, and T_i is the set of tokens in that sequence. Percentiles are computed over signed relevance scores for the target class, including negative scores (evidence for the opposite class). This yields a normalised score in $[0, 100]$ indicating how highly the token ranks relative to all other tokens in the sequence. The final metric is obtained by averaging these percentile scores across all token occurrences, allowing multiple tokens of interest to contribute within the same sequence.

F.5 Sizes of Attack Dataset

Table 28 provides information about the sizes of the attack dataset.

Approach	Model	Dataset	Expl. Method	Train.	Val.	Train. Frac.	Val Frac.
Overall Top Tokens	BERT	SST-2	LRP	44,595	726	0.66	0.83
	BERT	SST-2	GAE	22,473	475	0.33	0.54
	ALBERT	SST-2	LRP	67,349	872	1.00	1.00
	ALBERT	SST-2	GAE	67,349	872	1.00	1.00
	BERT	IMDb	LRP	24,966	12,486	1.00	1.00
	BERT	IMDb	GAE	22,831	11,345	0.91	0.91
	ALBERT	IMDb	LRP	25,000	12,500	1.00	1.00
	ALBERT	IMDb	GAE	25,000	12,500	1.00	1.00
Token Elevation	BERT	SST-2	LRP	47,243	866	0.70	0.99
	BERT	SST-2	GAE	67,349	872	1.00	1.00
	ALBERT	SST-2	LRP	30,666	661	0.46	0.76
	ALBERT	SST-2	GAE	67,349	872	1.00	1.00
	BERT	IMDb	LRP	24,974	12,488	1.00	1.00
	BERT	IMDb	GAE	24,464	12,204	0.98	0.98
	ALBERT	IMDb	LRP	16,674	8,256	0.67	0.66
	ALBERT	IMDb	GAE	18,015	9,062	0.72	0.72

Table 28: D_{fool} training and validation set coverage by approach, model, dataset, and explanation method (target tokens differ). The table shows wide coverage for almost all experiments.