

RAZOR: Sharpening Knowledge by Cutting Bias with Unsupervised Text Rewriting

Shuo Yang*, Bardh Prenkaj*, Gjergji Kasneci

Technical University of Munich
name.surname@tum.de

Abstract

Despite the widespread use of LLMs due to their superior performance in various tasks, their high computational costs often lead potential users to opt for the pretraining-finetuning pipeline. However, biases prevalent in manually constructed datasets can introduce spurious correlations between tokens and labels, creating so-called *shortcuts* and hindering the generalizability of fine-tuned models. Existing debiasing methods often rely on prior knowledge of specific dataset biases, which is challenging to acquire a priori. We propose RAZOR (Rewriting And Zero-bias Optimization Refinement), a novel, unsupervised, and data-focused debiasing approach based on text rewriting for shortcut mitigation. RAZOR leverages LLMs to iteratively rewrite potentially biased text segments by replacing them with heuristically selected alternatives in a shortcut space defined by token statistics and positional information. This process aims to align surface-level text features more closely with diverse label distributions, thereby promoting the learning of genuine linguistic patterns. Compared with unsupervised SoTA models, RAZOR improves by 3.5% on the FEVER and 6.5% on MNLI and SNLI datasets according to the F1 score. Additionally, RAZOR effectively mitigates specific known biases, reducing bias-related terms by $\times 2$ without requiring prior bias information, a result that is on par with SoTA models that leverage prior information. Our work prioritizes data manipulation over architectural modifications, emphasizing the pivotal role of data quality in enhancing model performance and fairness. This research contributes to developing more robust evaluation benchmarks for debiasing methods by incorporating metrics for bias reduction and overall model efficacy.

Code — <https://github.com/ShuoYangtum/RAZOR>

Appendices — <https://arxiv.org/abs/2412.07675>

1 Introduction

With the progression of instructional learning, prominent LLMs such as ChatGPT (Ouyang et al. 2024) and Vicuna (Zheng et al. 2023) are extensively utilized across diverse domains owing to their versatility and outstanding performance (Chang et al. 2024). However, their utility is significantly constrained by the high training and API licensing costs. As a result, numerous researchers in domains like

*These authors contributed equally.

Spielberg is a great spinner of a yarn, however this time he just didn't do it for me. (f_θ : ✓, Φ : ✗)

The benefits of a **New York Subway** system are that a person can get from A to B without being stuck in traffic and subway trains are faster than buses. (f_θ : ✗, Φ : ✓)

Figure 1: Example of spurious correlation in sentiment classification tasks, where a classifier f_θ takes *Spielberg* and *New York Subway* as shortcuts and makes wrong predictions w.r.t. the ground truth (Φ). The classifier concentrates on the bold tokens to make the prediction; however, the underlined tokens might be more useful in producing the correct label.

healthcare, education, and e-commerce adhere to the conventional approach of using open-source language models followed by fine-tuning (Strubell, Ganesh, and McCallum 2019).

LLMs absorb linguistic features from extensive corpora during pre-training by predicting specific lexical units. This foundational learning phase is then followed by fine-tuning domain-specific data, enhancing their ability to perform domain-specific tasks effectively.

Unlike linguistic features, surface biases are reported as commonly existing in manually curated datasets and can detrimentally affect the fine-tuning phase (Jiménez, Mendes Oliveira, and Soares 2020). In other words, models may achieve untruthfully elevated performance during training by exploiting specific shortcuts. For instance, analyses of the FEVER dataset (Thorne et al. 2018), used for fact verification tasks, revealed a pronounced association between the occurrence of negations within the claim and the corresponding label “refutes” (Karimi Mahabadi, Belinkov, and Henderson 2020). This suggests that models might erroneously generalize that claims containing negations are invariably fake.

The contemporary debiasing methods are relatively effective but require prior knowledge of biases. For example, recent approaches like CLEVER (Xu et al. 2023a) stand on the prior knowledge that shortcuts originate solely from tokens of specific parts within the input. Therefore,

they can estimate and correct the impact of shortcuts. However, such prior knowledge is scarce and difficult to obtain. Furthermore, bias types vary significantly with changes in datasets (Geirhos et al. 2020), severely limiting the generality of contemporary methods based on supervised learning.

As an unsupervised improvement strategy, our rationale lies in the limitations of LLMs to explicitly identify specific tokens that serve as shortcuts (Thomas McCoy, Pavlick, and Linzen 2020). Consequently, we propose equalizing all token-related surface¹ features across samples from the different classes. By achieving balance among these features, models are prevented from leveraging their potentially spurious correlations with labels. This approach thereby compels deeper semantic learning and augments overall robustness.

The main contributions of this paper are as follows:

- We introduce RAZOR – short for **R**ewriting **A**nd **Z**ero-bias **O**ptimization **R**efinement – a novel unsupervised debiasing technique that mitigates shortcuts in NLP models through iterative text rewriting.
- To the best of our knowledge, we are the first to formalize the definition of shortcuts about an underlying classifier and the ground truth of the sentences. Our theoretical findings help us understand RAZOR’s effectiveness in mitigating spurious correlations in real-world datasets without prior knowledge.
- RAZOR improves over the SoTA, achieving a 3.5% performance increase on the FEVER dataset and a 6.5% boost on MNLI and SNLI datasets.
- RAZOR reduces bias-related terms by $\times 2$, matching the performance of SoTA models without requiring prior knowledge of bias. This makes RAZOR a crucial advancement for tackling bias in NLP, particularly in scenarios where explicit bias information is inaccessible.
- Our approach emphasizes the understanding of shortcuts and highlights the power of data-driven techniques over architectural changes in bias mitigation. This reinforces the critical role of data quality in enhancing model fairness and performance while establishing a foundation for developing stronger evaluation benchmarks in future debiasing research.

2 Related Work

In recent years, significant research efforts have been directed toward identifying and mitigating spurious correlations in NLP tasks. Our work strictly relates to *creating balanced and less biased datasets* and *counterfactual data generation* lines of research.

Creating Balanced and Less Biased Datasets. Wu et al. (2022) present a data generation approach aimed at mitigating spurious correlations in natural language inference datasets through diverse data generation. The authors fine-tune GPT-2 as their data generation model to sequentially engender an instance’s premise, label, and hypothesis. Additionally, they apply unlikelihood training to mitigate label inconsistency phenomena, improving the label quality at

¹Surface features do not contain any semantic/linguistic characteristics in a sentence.

generation time. Finally, they add a consistency step (Bar-tolo et al. 2021; Lewis et al. 2021) – i.e., using the model’s confidence for filtering purposes – to improve the quality of the generated dataset further. After this process, the authors reject samples that contribute to the high spurious correlations between task-independent features of the samples and their labels via z-filtering and produce a debiased dataset. Bras et al. (2020) propose an iterative greedy algorithm that adversarially filters out spurious biases from data by relying on AFLite (Sakaguchi et al. 2020). Interestingly, the authors aim to mitigate performance leakage when datasets present spurious biases.² CrossAug (Lee et al. 2021) reduces dataset bias by creating contrastive pairs in fact-checking tasks. It uses a two-stage pipeline where the first stage generates negative claims by altering positive claims, and the second stage modifies the evidence to support these negative claims. In this way, CrossAug enhances a classification model’s ability to reason more accurately based on the provided evidence piece of text. Finally, CLEVER (Xu et al. 2023a) leverages counterfactual inference for debiasing in fact-checking tasks. Instead of relying on data augmentation or training-phase adjustments, CLEVER focuses on the inference stage. It does so by training two models, a claim-evidence fusion model, and a claim-only model, and then subtracting the output of the claim-only model from the fusion model to isolate and remove biases. EDA (Wei and Zou 2019) involves four simple techniques: synonym replacement, random insertion, random swap, and random deletion, all modifying the text while preserving its original meaning. These operations generate diverse training examples, helping models generalize better by exposing them to varied linguistic patterns. The approach is designed to be easy to implement and computationally inexpensive. EDA mitigates overfitting and improves model robustness by focusing on random, controlled transformations. Schuster et al. (2019) demonstrate that models trained solely on claims (without evidence) can achieve high accuracy due to dataset biases, particularly the presence of giveaway phrases that correlate with specific labels. To mitigate this issue, they create a “Symmetric Test Set”, which eliminates these biases by manually generating balanced claim-evidence pairs. Additionally, they propose a regularization method – namely ReW – that reweights training examples to reduce the influence of these biases during model training. Mahabadi, Belinkov, and Henderson (2020) propose “Product of Experts” – namely PoE – which combines predictions from a bias-only model and the main model, focusing the training on examples that are not biased.

Differently from the above works, we defer the rewriting process to GPT-3.5-Turbo, which has demonstrated advantages to mere heuristic-based methods that change synonyms (Wei and Zou 2019), negate verbs to engender new sentences (Lee et al. 2021), or reweigh the sentences to downgrade biases (Schuster et al. 2019). Similarly to (Wu et al. 2022), we use a consistency mechanism. However, we give GPT-3.5-Turbo the original evidence, the newly gener-

²Spurious biases in datasets make benchmarks artificially easier as models learn to overly rely on these biases instead of learning to generalize. This also hurts out-of-domain generalization.

ated claim, and ask it to tell whether it supports/refutes the new claim and repeat the generation process until the class distributions are aligned (see Sec. 3.3).

Counterfactual Data Generation. Lately, augmenting models with counterfactual data has become a favored method for reducing spurious correlations and enhancing model robustness. Kaushik and Lipton (2018) initially employed human workers to generate counterfactual examples for augmentation. Their findings indicate that counterfactually augmented data effectively mitigates spurious patterns in training data. However, this approach is costly, time-intensive, and often results in simple perturbations. Wu et al. (2021) and Ross et al. (2022) leverage text generation models to produce counterfactual data. These models require fine-tuning with specific perturbation types and have various limitations, including the untargeted and unlabelled generation process and the restricted perturbation types. To introduce new perturbation types, the models need retraining. Unlike the previous two methods, Chen et al. (2023) develop DISCO, a framework for distilling counterfactuals with large language models to help distinguish between spurious and genuine correlations. DISCO does prompt engineering to generate phrasal perturbations with a large general language model. Then, a task-specific teacher model filters these generations to distill high-quality counterfactual data to enrich the original training set. Lastly, Xu et al. (2023b) focus on counterfactual debiasing for fact verification, generating counterfactual scenarios to identify and mitigate spurious correlations in verification tasks.

3 Method

3.1 Problem Formulation

We consider binary classifications of short text documents, e.g., sentiment classification of reviews. Let us assume that we have a set of i.i.d. labelled samples – i.e., dataset – $\mathcal{D} = \{d_1, \dots, d_n\}$. Each sample in \mathcal{D} is a document, i.e., a sequence of tokens, $d_i = \{t_1, t_2, \dots, t_m\}$ with a corresponding label $y_i \in \mathcal{Y}$ where w.l.o.g. $\mathcal{Y} = \{0, 1\}$. To classify a document d_i , we first transform it into a feature vector \mathbf{x}_i via a designated transformation function $g : \mathcal{D} \rightarrow \mathcal{X} \subseteq \mathbb{R}^u$, where \mathcal{X} is the domain of features. Then, \mathbf{x}_i is classified by a classification function $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ with model parameters θ . The parameters θ are typically estimated from \mathcal{D} by optimizing a certain loss function $\theta^* \leftarrow \arg \min_\theta \mathcal{L}(\mathcal{D}, \theta)$.

Generally, the classification function f_θ adopted in NLP is a deep learning model (see among others (Naseem et al. 2020; Li et al. 2024a)). Alas, as noticed in (Izmailov et al. 2022), deep classifiers are known to rely on spurious features – i.e., patterns that are correlated with the target on the training data but not inherently relevant to the learning problem, such as the image backgrounds when classifying the foreground. In NLP problems, such as sentiment classification tasks, certain tokens (e.g., words) can be associated with positive/negative classes due to their co-occurrence with the labels in the training data. Inspired by (Wang et al. 2022), in Fig. 1, we show how the classifier exploits the shortcut terms *Spielberg* and *New York Subway* to misclassify the sentences. While shortcut mitigation has been extensively

studied, to the best of our knowledge, there is a lack of a uniform definition of what a shortcut represents. Hence, we provide the reader with a formal definition of shortcut tokens. Notice that each document d_i gets tokenized into a sequence of tokens $\{t_1, t_2, \dots, t_m\}$, e.g., at a word-level, character-level, or sub-word level before being transformed into a feature vector x_i . Lastly, assume that $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$ represents the ground truth function.

Definition 1 Given a document of tokens $d_i = \{t_1, t_2, \dots, t_m\} \in \mathcal{D}$, a feature transformation function $g : \mathcal{D} \rightarrow \mathcal{X} \subseteq \mathbb{R}^u$, a classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, and a ground truth function $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$, a set of tokens $\hat{d}_i \subset d_i$ is a shortcut if the following conditions are satisfied:

$$f_\theta(g(\hat{d}_i)) = f_\theta(g(d_i)) \quad (1)$$

$$f_\theta(g(d_i)) \neq \Phi(g(d_i)) \quad (2)$$

$$|\hat{d}_i| \leq |d_i \setminus \hat{d}_i| \quad (3)$$

For the sake of the argument, let us assume that f_θ is a transformer. We indicate with $h : \mathcal{X} \rightarrow \mathbb{R}^\ell$ a score-attribution function to each input feature. Notice that the following derivation can be applied for any attribution function – either post-hoc, such as SHAP (Lundberg and Lee 2017) or not. We leave the derivation under a generic attribution function for future work and focus on the attention mechanism of a transformer in this work.

Attention mechanism. Let $h : \mathcal{X} \rightarrow \mathbb{R}^\ell$ be the attention mechanism that, given in input a document $d_i = \{t_1, t_2, \dots, t_m\}$, produces attention weights \mathbf{z}_j for each token $t_j \in d_i$. In this work, we assume that $h(g(t_j)) = \mathbf{z}_j \forall t_j \in d_i$.

Essential vs. non-essential tokens. According to Definition 1, we have two sets of tokens for each document d_i : i.e., the set of “essential” tokens \hat{d}_i , and that of “non-essential” ones $d_i \setminus \hat{d}_i$. The essential tokens are the ones that contribute the most (i.e., their attention scores are high) to the prediction. In contrast, non-essential tokens do not significantly contribute to the prediction (Vaswani et al. 2017). Because the prediction does not change if we use the entire set of tokens in the document d_i , intuitively, one can conclude that the attention scores of the set of non-essential are lower than those of the essential set of tokens (see Lemma 1).

Lemma 1 Given a document of tokens $d_i = \{t_1, t_2, \dots, t_m\} \in \mathcal{D}$, and an attention scores attribution function $h : \mathcal{X} \rightarrow \mathbb{R}^\ell$, if $\hat{d}_i \subset d_i$ is a shortcut, then

$$\frac{H(\hat{d}_i)}{|\hat{d}_i|} \geq \frac{H(d_i \setminus \hat{d}_i)}{|d_i \setminus \hat{d}_i|} \quad (4)$$

where

$$H(\hat{d}_i) = \left\| \sum_{t_j \in \hat{d}_i} h(g(t_j)) \right\|_2 \quad (5)$$

$$H(d_i \setminus \hat{d}_i) = \left\| \sum_{t_j \in d_i \setminus \hat{d}_i} h(g(t_j)) \right\|_2 \quad (6)$$

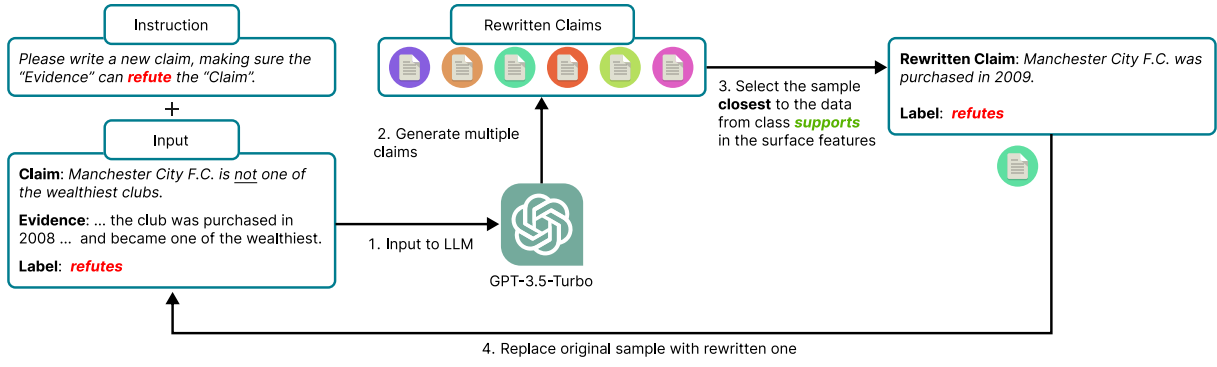


Figure 2: An example of RAZOR’s application in the fact-checking task. The task aims to determine whether a piece of evidence from Wikipedia **supports** or **refutes** a claim. The instance here is sampled from the FEVER dataset, where a negation word “not” has been reported to exhibit a spurious correlation with the class **refutes**.

We invite the reader to check Appendix A for the proof of Lemma 1. In practice, we observe that shortcuts are much smaller than the rest of the tokens in a sentence, which paves the way for efficient methods to identify the subset of tokens that represent shortcuts.

3.2 Shortcut Identification

Throughout the rest of the paper, we use document and sentence interchangeably. In this paper, we propose **R**ewriting **A**nd **Z**ero-bias **O**ptimization **R**efinement – namely RAZOR – a sentence rewriting strategy that mitigates shortcuts and promotes dataset debiasing. RAZOR aims to rewrite \mathcal{D} by balancing the number of essential words between positive and negative sentences. However, the cost of rewriting all sentences in \mathcal{D} is exceedingly high. As a heuristic optimization approach, we first filter a subset $\hat{\mathcal{D}}$ that most likely includes shortcuts – according to Eq. 12 – rewriting it and repeating this filter-rewriting process.

To find $\hat{\mathcal{D}}$, we define a surface feature extraction function $\hat{g} : \mathcal{D} \rightarrow \mathbb{R}^\lambda$ which maps all samples into a shortcut space. Notably, the specific types of shortcuts are strongly task-dependent. In this work, without any prior knowledge, we provide a general identification pattern with the two essential factors by following Definition 1: i.e., (1) the **significance** of tokens, and (2) the (relative) **positions** of tokens.

We first use the TF-IDF score to calculate the importance of a specific token $t_j \in d_i$. To achieve this, we define $S : \mathcal{D} \rightarrow \mathbb{R}$ as in Eq. (7):

$$S(t_j) = \frac{n(t_j, d_i)}{|d_i|} \log \frac{|\mathcal{D}|}{|\{d_k \mid \exists d_k \in \mathcal{D} \text{ s.t. } t_j \in d_k\}|} \quad (7)$$

where $n(t_j, d_i)$ is the number of occurrences of token t_j in d_i . Applying Eq. (7) offers a key advantage: i.e., the TF-IDF-based measurement keeps the balance of token frequency, while word frequency is one of the factors most likely to cause shortcuts (Thomas McCoy, Pavlick, and Linzen 2020). For example, suppose a token holds considerably different frequencies in positive versus harmful data. As per how transformers work, this essential token will hold higher attention scores during training, leading the model f_θ to focus more on its occurrence rather than the semantic information of the sample – i.e., this token is considered a shortcut.

Second, we inject positional information of words through the positional embeddings since the model f_θ may learn shortcuts from the relative positions of tokens. Specifically, we compare positional embeddings of the same token from samples with different labels to ensure they are closely aligned. Without prior knowledge, we use a fixed positional encoding function $\tau : \mathbb{N} \rightarrow \mathbb{R}^\lambda$ based on sine and cosine functions

$$\tau(\text{pos}) = \left[\text{PE}(\text{pos}, 0), \dots, \text{PE}(\text{pos}, \lambda - 1) \right], \quad (8)$$

with

$$\text{PE}(\text{pos}, k) = \begin{cases} \sin\left(\frac{\text{pos}}{10000^{2k/\lambda}}\right) & k \equiv 0 \pmod{2} \\ \cos\left(\frac{\text{pos}}{10000^{2k/\lambda}}\right) & k \equiv 1 \pmod{2} \end{cases} \quad (9)$$

where pos is the token’s position in a target document, k is the dimension index in the embedding vector, and λ is the dimensionality of the positional embedding vector. Afterwards, we compute the mean of the products of S for all tokens in a document d_i as its shortcut feature mapping

$$\hat{g}(d_i) = \frac{\sum_{t_j \in d_i} S(t_j) * \tau(j)}{|d_i| - 1}, \quad (10)$$

assuming that t_j occurs in position j in the document d_i , and $*$ is the scalar product. Finally, let us use $\vartheta(d_i, \Phi)$ to indicate

$$\vartheta(d_i, \Phi) = \{d_j \mid d_j \in \mathcal{D} \text{ s.t. } \Phi(d_i) \neq \Phi(d_j)\}. \quad (11)$$

In other words, given a document $d_i \in \mathcal{D}$, and a ground truth function Φ , as introduced in Definition 1, $\vartheta(d_i, \Phi)$ produces a set of documents whose label is different with $\Phi(d_i)$. In this way, we define a *shortcut score* $\gamma : \mathcal{D} \rightarrow \mathbb{R}$ as in Eq. 12.

$$\gamma(d_i) = 1 - \frac{1}{|\vartheta(d_i, \Phi)|} \sum_{d_j \in \vartheta(d_i, \Phi)} \frac{\hat{g}(d_i) \cdot \hat{g}(d_j)}{\underbrace{\|\hat{g}(d_i)\| \|\hat{g}(d_j)\|}_{\cos(\hat{g}(d_i), \hat{g}(d_j))}} \quad (12)$$

Note that the higher the shortcut score, the greater the token position distribution differs between the sample d_i , assuming it is positive, and the negative samples in $\vartheta(d_i, \Phi)$,

indicating that the model f_θ is more likely to learn shortcuts from d_i . We select the top documents with the highest shortcut values to form $\widehat{\mathcal{D}}$ and use the documents therein for rewriting.

3.3 Rewriting and Filtering

After obtaining the set of shortcut candidates, we rewrite each sentence $d_i \in \widehat{\mathcal{D}}$ into several candidate sentences maintaining the original label $\Phi(d_i)$. We then substitute d_i with the rewritten sentence least likely to contain shortcuts. In Fig. 2, we show how this rewriting process works in a fact-checking task. Here, we aim for the LLM to generate a new *claim* based on the same *evidence* from Wikipedia while ensuring that the evidence still supports (or refutes) the newly generated claim. To rewrite the sentence d_i , we construct a prompt for an LLM that generates rewritten candidates. The prompt contains two components: i.e., an *instruction*, which is a general explanation of the rewriting task for d_i including $\Phi(d_i)$; and the *input* d_i itself. We ask the LLM to engender rewritten sentences while maintaining the original label $\Phi(d_i)$. Hence, we rely on a self-consistency mechanism to ensure the quality of generated sentences. Specifically, in the fact-checking scenario, we provide the LLM with the original evidence and the newly generated claim and query it to tell us whether the evidence supports/refutes the new claim. More formally, let $G_\alpha : \mathcal{D} \times \mathcal{Y} \rightarrow \mathcal{D}$ be the LLM we are querying to generate new sentences respecting the original label. Additionally, let $G_\beta : \mathcal{D} \times \mathcal{Y} \rightarrow \{0, 1\}$ be another LLM³ that verifies whether the generated sentence has the original label. More formally, we accept an engendered sentence $d_i^* = G_\alpha(d_i, \Phi(d_i))$ if Eq. (13) is satisfied.

$$\mathbb{I} \left[G_\beta \left(G_\alpha(d_i, \Phi(d_i)), \Phi(d_i) \right) = 1 \right] \quad (13)$$

Once, the rewriting process is completed for a particular d_i , we obtain a set of generated documents $\varphi(d_i) = \{d_i^* \mid d_i^* = G_\alpha(d_i, \Phi(d_i))\}$. From $\varphi(d_i)$, we choose the one scoring the least according to Eq. 12 and insert it in the original dataset \mathcal{D} . The original document d_i is removed from the “rewritten” dataset as we aim to mitigate the shortcuts that it might have contained. In this way, RAZOR maintains the same size as the original dataset throughout the rewriting process. We repeat this process until the cosine similarity between the embeddings in the surface feature space of the documents of each label is maximal (see Eq. 14).

$$\begin{aligned} & \text{maximize} \sum_{d_i \in \mathcal{D}_0} \sum_{d_j \in \mathcal{D}_1} \cos(\hat{g}(d_i), \hat{g}(d_j)) \\ & \text{s.t. } \mathcal{D}_y = \{d \mid d \in \mathcal{D} \wedge \Phi(d) = y\} \end{aligned} \quad (14)$$

4 Experiments

Datasets. Fact-checking tasks aim to determine the veracity of a claim from social media based on existing facts.

³Note that we exploit these LLMs in zero-shot. Therefore, G_α and G_β can be the same kind of LLM – e.g., GPT-4o – as long as there are two different sessions between the generation and the verification processes.

Here, we use the FEVER dataset (Thorne et al. 2018), which consists of claims and evidence retrieved from Wikipedia. The training objective is to classify whether the evidence supports or refutes the claim. Due to the prevalence of generating non-factual claims by simply negating factual claims, a significant negation bias has been reported in the literature.

Natural Language Inference (NLI) aims to determine the semantic relationship between a pair of sentences. Given a premise and a hypothesis, the goal is to classify the relationship as entailment, contradiction, or neutral. Here, we rely on the Multi-Genre Natural Language Inference (MNLI) (Williams, Nangia, and Bowman 2018) corpus and the Stanford Natural Language Inference (SNLI) (Bowman et al. 2015) corpus, which provides a large number of sentence pairs annotated with their logical relationship.

Experimental setup. Our preprocessing steps involve truncating sequences to 512 tokens. We rely on three classification models to measure RAZOR’s effectiveness, namely BERT, RoBERTa, and DistilBERT – taken from HuggingFace⁴ – followed by a linear classification layer. We train the classifiers using the LoRa algorithm (Hu et al. 2022) and the AdamW optimizer with a batch size of 16 and a learning rate of 3×10^{-5} . In the sentence rewriting, we set $G_\alpha = \text{GPT-3.5-Turbo}$ and $G_\beta = \text{GPT-3.5-Turbo}$. Here, we control the diversity of generated sequences by setting the top-p value to 0.9 and the temperature to 0.7. To provide details on the impact of the selected LLM to rewrite the sentences and verify their labels, we change both G_α and G_β to LLaMA-3.1-8B-Instruct (see Tables 1 and 2).

4.1 Discussion

RAZOR is a more robust dataset debiaser than SoTA methods on fact-checking tasks with, respectively, $\sim 2\%$ and $\sim 3.63\%$ increase in Accuracy and F1 over the second-best. We train BERT, RoBERTa, and DistilBERT classifiers on the original FEVER dataset to represent our baseline (hereafter Original). Then, for each SoTA method, we train the two classifiers on the modified version of FEVER according to the proposed algorithms. Notice that RAZOR and other SoTA methods only rewrite the training set of the dataset. The test set remains unchanged. We follow this strategy to assess the impact of the rewriting strategy w.r.t. the test set’s original peculiarities, which might still contain shortcuts. We expect that if the rewriting strategy truly alleviates the shortcut problem in the training set, then in the test set, we will see an increase in performance since the false negative/positive ratio should decrease⁵. Table 1 shows this effect on the two classifiers mentioned above for RAZOR and SoTA methods. Notice how, although second-best after CrossAug, RAZOR effectively removes the shortcut effect on both classifiers w.r.t. the original baseline. We acknowledge that RAZOR is behind CrossAug in

⁴For BERT, we use the `bert-base-uncased` pre-trained model. For RoBERTa, see <https://huggingface.co/FacebookAI/roberta-base>. For DistilBERT, see <https://huggingface.co/distilbert/distilbert-base-uncased>.

⁵As per Definition 1, a shortcut causes classifiers to have false positives/negatives.

		FEVER				FEVER-Adversarial			
		BERT		RoBERTa		BERT		RoBERTa	
		Accuracy	F1	Accuracy	F1	Accuracy	F1	Accuracy	F1
Rewriting Strategy	Original	87.93±0.51	86.65±0.78	90.05±0.48	88.90±0.69	64.10±1.21	69.81±1.45	66.99±1.30	73.13±1.82
	EDA	87.99±0.30	87.03±0.50	89.25±0.32	89.00±0.66	65.20±0.99	69.55±0.86	66.35±1.54	69.98±1.97
	CrossAug	90.28±0.42	89.37±0.54	91.00±0.46	91.98±0.78	63.58±1.20	66.74±1.56	67.02±1.44	72.04±1.77
	ReW	87.50±0.88	87.76±1.19	89.09±1.21	89.44±2.01	64.52±2.21	69.85±2.30	67.23±2.03	74.20±2.68
	PoE	86.25±1.78	86.98±1.89	90.20±1.78	89.79±1.90	65.20±2.76	70.45±2.58	68.00±2.49	72.12±2.61
	CLEVER	85.13±1.20	85.33±2.03	86.55±1.72	84.54±2.33	56.16±2.54	58.93±2.27	62.63±2.96	63.35±2.35
	RAZOR (w/ LLaMA)	87.97±1.20	87.22±2.03	90.59±2.32	89.11±2.56	66.03±1.95	69.61±1.88	68.77±3.01	73.76±2.50
	RAZOR (w/ GPT)	88.73±2.23	88.39±2.20	90.45±1.99	90.22±2.11	66.66±1.98	73.78±1.37	69.21±1.64	76.08±2.02

Table 1: Comparison of RAZOR with SoTA methods on the FEVER and FEVER-Adversarial datasets – the datasets have the same training set but different test sets. We show the performance of the classifiers in terms of accuracy, and F1 scores after the shortcuts within the dataset have been modified using each debiasing method, including RAZOR (ours). *Original* means that the original training set has not been modified, and we report the performances of the two classifiers as is.

		MNLI (MNLI)		SNLI (MNLI)		MNLI (SNLI)		SNLI (SNLI)	
		Accuracy	F1	Accuracy	F1	Accuracy	F1	Accuracy	F1
BERT		86.70±1.21	86.85±1.79	83.05±1.38	84.41±2.01	70.15±2.89	62.75±2.77	91.65±0.80	91.74±0.91
w/ RAZOR (GPT)		88.50±1.32	87.90±1.80	84.75±1.72	85.07±2.44	74.10±3.20	67.86±3.57	93.20±1.11	93.16±1.26
w/ RAZOR (LLaMA)		86.80±1.32	86.88±1.80	84.10±1.52	83.51±2.00	74.25±2.92	68.73±3.14	92.90±1.27	92.81±1.32
Δ (GPT)		+2.08	+1.44	+2.04	+0.78	+5.63	+8.14	+1.69	+1.55
Δ (LLaMA)		+0.12	+0.04	+1.26	-1.06	+5.85	+9.56	+1.36	+1.17
RoBERTa		90.85±1.10	90.75±1.53	90.30±1.27	91.08±1.55	85.35±2.04	84.93±1.48	95.05±0.72	95.01±0.80
w/ RAZOR (GPT)		91.15±1.26	91.14±1.48	91.45±1.29	91.42±1.68	86.50±2.20	86.22±1.72	95.15±0.70	95.17±0.98
w/ RAZOR (LLaMA)		91.55±1.38	91.45±1.70	90.90±1.29	91.13±1.68	85.50±1.98	85.76±2.24	95.60±0.70	95.71±0.98
Δ (GPT)		+0.33	+0.42	+1.27	+0.37	+1.34	+1.52	+0.11	+0.17
Δ (LLaMA)		+0.77	+0.77	+0.66	+0.05	+0.18	+0.98	+0.58	+0.74
DistilBERT		83.50±1.23	82.97±1.77	77.23±1.85	77.38±1.86	65.55±2.30	69.90±1.78	88.82±1.22	87.78±1.67
w/ RAZOR (GPT)		85.14±1.66	85.20±1.84	80.22±1.90	78.82±2.21	68.58±2.16	72.80±1.83	91.82±0.88	91.80±1.03
w/ RAZOR (LLaMA)		84.00±1.35	83.98±1.78	77.35±1.43	77.62±1.99	66.33±2.00	71.30±1.85	90.30±1.26	90.23±1.97
Δ (GPT)		+1.96	+2.69	+3.87	+1.86	+4.62	+4.14	+3.38	+4.58
Δ (LLaMA)		+0.60	+1.22	+0.16	+0.31	+1.19	+2.00	+1.67	+2.79

Table 2: RAZOR’s impact on the classifiers on the NLI datasets in terms of Accuracy and F1 score. Note that we assess the performances of the classifiers across different train-test set combinations – e.g., training on MNLI and testing on SNLI – to mitigate the fact that the original test set contains similar shortcuts as the original training set. Within parentheses, we indicate the training set; outside of them, we indicate the test set.

terms of performance because the latter enriches the dataset with contrastive – i.e., positive-negative – claim-evidence pairs, where the evidence is rewritten differently from RAZOR, which only operates at the claim level. However, we show how RAZOR is more robust than CrossAug in the FEVER-Adversarial dataset.⁶ The sentences in the test set of FEVER-Adversarial do not contain shortcuts. Hence, we use the adversarial test set to demonstrate RAZOR’s debiasing abilities. Notice how CrossAug’s performance is systematically lower than that of RAZOR. Because CrossAug negates the claim and evidence to generate other factual instances, it might happen that this newly generated claim-evidence pair can itself contain shortcuts, exacerbating the classifier’s reasoning capabilities. However, since FEVER’s test set contains shortcuts, CrossAug’s original performances are not faithful to assess the goodness of this debiasing strategy. In particular, this is evident in FEVER-Adversarial, where simple negations of claim-evidence pairs are not optimal. Contrarily, RAZOR reduces shortcuts’ effect and achieves consistently better results than all SoTA methods.

⁶The training set of FEVER and FEVER-Adversarial is the same. The test sets are different.

RAZOR demonstrates cross-dataset generalizability in multi-class classification tasks. To validate the generalizability of our method, we rely on MNLI and SNLI for a multi-class classification test. Similarly to the above process, we fine-tuned the classifiers as baselines to determine whether the logical relationship between two sentences is entailment, contradiction, or neutral. Subsequently, we used RAZOR to rewrite the training sets and achieve improvements in Accuracy and F1 scores for the two classifiers. With the lack of corresponding attacking datasets like FEVER-Adversarial, we exchanged MNLI and SNLI test sets to simulate a real-world data distribution.

Our experiments demonstrate that even widely used datasets exhibit shortcut problems. We observed that by applying RAZOR, the baseline classifiers achieved improvements of approximately 0.1 to 8 points across all datasets and metrics. The most significant improvement is observed in the SNLI-MNLI training-test set combination. This suggests that SNLI may contain more shortcuts than MNLI, resulting in a performance difference of over 20 points for BERT across different test sets. However, after applying RAZOR, a substantial portion of the shortcuts in the data was

removed, leading to an increase of 8.14 points in F1 score and 5.63 points in accuracy for BERT. Moreover, we find that RAZOR performs better on all classifiers overall. This is because models with better overall performances are less susceptible to shortcut issues. Therefore, RAZOR demonstrates significant potential for debiasing smaller models.

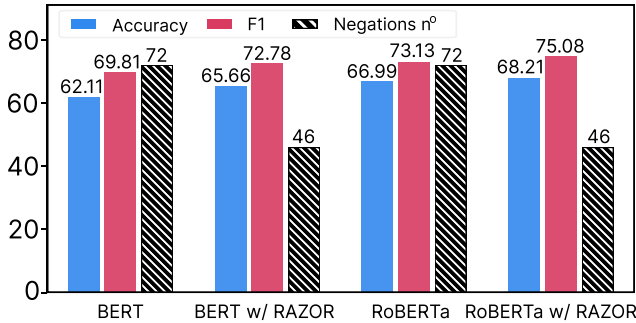


Figure 3: Effect of shortcut-related terms for BERT and RoBERTa with and without RAZOR on 500 randomly sampled original-rewritten pairs on the FEVER dataset. We then test on the FEVER-Adversarial set.

RAZOR effectively reduces the number of negations considered shortcuts in FEVER by 36.2%. To further quantify the effectiveness of our approach in mitigating shortcuts, we perform a statistical analysis of shortcut-related terms, specifically negations, within both the original and rewritten FEVER training set. Following the methodology presented in (Li et al. 2024b), we train models with randomly sample 500 original-rewritten pairs from the two sets, count the occurrences of the words "no" and "not" in their claims, and repeat the experiments detailed in Table 1 for the test set of the FEVER-Adversarial dataset. We illustrate the results in Fig. 3. Notice that, by applying RAZOR, the occurrence of shortcut-related terms is reduced by approximately 40% without any prior knowledge of these terms. Furthermore, eliminating shortcuts significantly improved models' generalizability, leading to an increase of, respectively, ~ 3 and ~ 2 percentage points for BERT and RoBERTa in both Accuracy and F1 scores when we test on the FEVER-Adversarial set.

4.2 Ablation Studies

To ensure that the performance improvement of our model does not stem from GPT itself, we first ruled out potential information leakage by calculating repetition rates. Subsequently, we replaced our shortcut space mapping with the model's token embeddings and employed the same method to demonstrate the effectiveness of our defined shortcuts. We illustrate the results in Fig. 4. We observe consistent trends in the Accuracy and F1 score training curves. RAZOR outperforms the baseline and significantly surpasses the sentence embedding-based method on RoBERTa and BERT.

Furthermore, RAZOR exhibits a slightly slower convergence rate than the baseline but meets equilibrium earlier than the embedding-based method. We attribute this to RAZOR's ability to mitigate the influence of surface features,




Figure 4: The performance of models trained on FEVER when evaluated on the original FEVER test set.

encouraging the classifiers to move beyond specific token biases and invest more time in learning linguistic features, ultimately resulting in superior performance. In contrast, embedding-based methods rely on sentence representations that mix surface and linguistic features, failing to separate the two effectively, preventing the classifiers from avoiding shortcuts. Moreover, the proximity of embeddings may obscure the classification hyperplane, inherently complicating the task and leading to mediocre performance. Given that both methods employ GPT-3.5-Turbo but produce different outcomes, we can safely conclude that the overall RAZOR's framework, rather than solely GPT in the rewriting process, is a primary driver in its effectiveness. Appendix B illustrates the performance difference between RAZOR operating on BERT with surface features and embedding features.

5 Conclusion

RAZOR is an effective dataset debiasing method. It outperforms SoTA methods in addressing shortcut problems in fact-checking and natural language inference tasks. RAZOR maintains dataset integrity by generating new sentences that adhere to the original labels and strategically removing the original documents while significantly reducing shortcut-induced biases. This is achieved through an iterative rewriting process, optimized until the KL divergence between the embeddings in the surface feature space is minimized. Our experiments consistently show that RAZOR enhances the performance of classifiers such as BERT and RoBERTa, particularly in challenging adversarial scenarios where traditional methods like CrossAug fail to deliver competitive results. Notably, RAZOR's superior performance on models like BERT, even compared to more robust models like RoBERTa, underscores its potential as a powerful debiasing tool, particularly for smaller and more resource-constrained models. The implications of our work extend beyond the datasets and models tested, offering a promising approach to improving model reliability across a wide range of machine-learning applications. While RAZOR has shown great promise, future work could explore its application to other types of biases or in conjunction with additional debiasing strategies. We believe that RAZOR represents a significant step forward in the ongoing effort to create more fair and accurate machine learning systems, ultimately contributing to more trustworthy AI applications in the real world.

References

- Bartolo, M.; Thrush, T.; Jia, R.; Riedel, S.; Stenetorp, P.; and Kiela, D. 2021. Improving Question Answering Model Robustness with Synthetic Adversarial Data Generation. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8830–8848. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In Márquez, L.; Callison-Burch, C.; Su, J.; Pighin, D.; and Marton, Y., eds., *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, 632–642. The Association for Computational Linguistics.
- Bras, R. L.; Swayamdipta, S.; Bhagavatula, C.; Zellers, R.; Peters, M.; Sabharwal, A.; and Choi, Y. 2020. Adversarial Filters of Dataset Biases. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1078–1088. PMLR.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; Ye, W.; Zhang, Y.; Chang, Y.; Yu, P. S.; Yang, Q.; and Xie, X. 2024. A Survey on Evaluation of Large Language Models. *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Chen, Z.; Gao, Q.; Bosselut, A.; Sabharwal, A.; and Richardson, K. 2023. DISCO: Distilling Counterfactuals with Large Language Models. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5514–5528. Toronto, Canada: Association for Computational Linguistics.
- Geirhos, R.; Jacobsen, J.-H.; Michaelis, C.; Zemel, R. S.; Brendel, W.; Bethge, M.; and Wichmann, F. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2: 665 – 673.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Izmailov, P.; Kirichenko, P.; Gruver, N.; and Wilson, A. G. 2022. On feature learning in the presence of spurious correlations. *Advances in Neural Information Processing Systems*, 35: 38516–38532.
- Jiménez, L.; Mendes Oliveira, H.; and Soares, A. P. 2020. Surface features can deeply affect artificial grammar learning. *Consciousness and Cognition*, 80: 102919.
- Karimi Mahabadi, R.; Belinkov, Y.; and Henderson, J. 2020. End-to-End Bias Mitigation by Modelling Biases in Corpora. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8706–8716. Online: Association for Computational Linguistics.
- Kaushik, D.; and Lipton, Z. C. 2018. How Much Reading Does Reading Comprehension Require? A Critical Investigation of Popular Benchmarks. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 5010–5015. Brussels, Belgium: Association for Computational Linguistics.
- Lee, M.; Won, S.; Kim, J.; Lee, H.; Park, C.; and Jung, K. 2021. CrossAug: A Contrastive Data Augmentation Method for Debiasing Fact Verification Models. In Demartini, G.; Zuccon, G.; Culpepper, J. S.; Huang, Z.; and Tong, H., eds., *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, 3181–3185. ACM.
- Lewis, P.; Wu, Y.; Liu, L.; Minervini, P.; Küttler, H.; Piktus, A.; Stenetorp, P.; and Riedel, S. 2021. PAQ: 65 Million Probably-Asked Questions and What You Can Do With Them. *Transactions of the Association for Computational Linguistics*, 9: 1098–1115.
- Li, X.; Wang, B.; Wang, Y.; and Wang, M. 2024a. Graph-based text classification by contrastive learning with text-level graph augmentation. *ACM Transactions on Knowledge Discovery from Data*, 18(4): 1–21.
- Li, Y.; Xu, M.; Miao, X.; Zhou, S.; and Qian, T. 2024b. Prompting Large Language Models for Counterfactual Generation: An Empirical Study. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 13201–13221. Torino, Italia: ELRA and ICCL.
- Lundberg, S. M.; and Lee, S. 2017. A Unified Approach to Interpreting Model Predictions. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 4765–4774.
- Mahabadi, R. K.; Belinkov, Y.; and Henderson, J. 2020. End-to-End Bias Mitigation by Modelling Biases in Corpora. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J. R., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 8706–8716. Association for Computational Linguistics.
- Naseem, U.; Razzak, I.; Musial, K.; and Imran, M. 2020. Transformer based deep intelligent contextual embedding for twitter sentiment analysis. *Future Generation Computer Systems*, 113: 58–69.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2024. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing*

- Systems, NIPS '22. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781713871088.
- Ross, A.; Wu, T.; Peng, H.; Peters, M.; and Gardner, M. 2022. Tailor: Generating and Perturbing Text with Semantic Controls. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3194–3213. Dublin, Ireland: Association for Computational Linguistics.
- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2020. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 8732–8740. AAAI Press.
- Schuster, T.; Shah, D. J.; Yeo, Y. J. S.; Filizzola, D.; Santus, E.; and Barzilay, R. 2019. Towards Debiasing Fact Verification Models. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 3417–3423. Association for Computational Linguistics.
- Strubell, E.; Ganesh, A.; and McCallum, A. 2019. Energy and Policy Considerations for Deep Learning in NLP. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. Florence, Italy: Association for Computational Linguistics.
- Thomas McCoy, R.; Pavlick, E.; and Linzen, T. 2020. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, 3428–3448. Association for Computational Linguistics (ACL).
- Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; and Mittal, A. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *NAACL-HLT*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, T.; Sridhar, R.; Yang, D.; and Wang, X. 2022. Identifying and Mitigating Spurious Correlations for Improving Robustness in NLP Models. In Carpuat, M.; de Marneffe, M.; and Ruíz, I. V. M., eds., *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, 1719–1729. Association for Computational Linguistics.
- Wei, J. W.; and Zou, K. 2019. EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 6381–6387. Association for Computational Linguistics.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122. Association for Computational Linguistics.
- Wu, T.; Ribeiro, M. T.; Heer, J.; and Weld, D. 2021. Polyjuice: Generating Counterfactuals for Explaining, Evaluating, and Improving Models. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 6707–6723. Online: Association for Computational Linguistics.
- Wu, Y.; Gardner, M.; Stenetorp, P.; and Dasigi, P. 2022. Generating Data to Mitigate Spurious Correlations in Natural Language Inference Datasets. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2660–2676. Dublin, Ireland: Association for Computational Linguistics.
- Xu, W.; Liu, Q.; Wu, S.; and Wang, L. 2023a. Counterfactual Debiasing for Fact Verification. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6777–6789. Toronto, Canada: Association for Computational Linguistics.
- Xu, W.; Liu, Q.; Wu, S.; and Wang, L. 2023b. Counterfactual debiasing for fact verification. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6777–6789.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E. P.; Zhang, H.; Gonzalez, J.; and Stoica, I. 2023. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *ArXiv*, abs/2306.05685.