

# Language Models Get a Gender Makeover: Mitigating Gender Bias with Few-Shot Data Interventions

Himanshu Thakur Atishay Jain\* Praneetha Vaddamanu\*

Paul Pu Liang Louis-Philippe Morency

Carnegie Mellon University

{hthakur, atishayj, pvaddama, pli, liang, morency}@andrew.cmu.edu

## Abstract

*Caution: this paper contains potentially offensive or upsetting model outputs.*

Societal biases present in pre-trained large language models are a critical issue as these models have been shown to propagate biases in countless downstream applications, rendering them unfair towards specific groups of people. Since large-scale retraining of these models from scratch is both time and compute-expensive, a variety of approaches have been previously proposed that de-bias a pre-trained model. While the majority of current state-of-the-art debiasing methods focus on changes to the training regime, in this paper, we propose data intervention strategies as a powerful yet simple technique to reduce gender bias in pre-trained models. Specifically, we empirically show that by fine-tuning a pre-trained model on only 10 de-biased (intervened) training examples, the tendency to favor any gender is significantly reduced. Since our proposed method only needs a few training examples, our few-shot debiasing approach is highly feasible and practical. Through extensive experimentation, we show that our debiasing technique performs better than competitive state-of-the-art baselines with minimal loss in language modeling ability.

## 1 Introduction

Recently, there has been a surge of interest in pre-trained large language models (LLM) in natural language processing (NLP). It has been shown that the pre-training + finetuning of a model drastically improves its performance on downstream tasks as the knowledge captured by the pre-training on a large corpus is transferred to the downstream application when finetuning the model. However, this also leads to societal biases like gender bias that were implicitly learned by the pre-trained models being transferred to crucial downstream applications like job recommendation engines (Zhao et al., 2019;

Barocas et al., 2017; Kurita et al., 2019). Analyzing and mitigating bias without requiring significant re-training or compute resources is crucial to the widespread adoption of LLMs in downstream applications.

Previous work (Nadeem et al., 2021), (Nangia et al., 2020a), (Cer et al., 2018) has attempted to quantify bias, and others such as Ravfogel et al. (2020) and Liang et al. (2021) have attempted to remove it algorithmically from the models. Closer to our work are data-manipulative techniques such as Zmigrod et al. (2019) and Maudslay et al. (2019) that modify the dataset and further fine-tune the model. In this paper, we propose simple data intervention strategies and show that they can mitigate gender bias in pre-trained models with the help of few-shot fine-tuning. Moreover, taking inspiration from Schick et al. (2021), we find that by utilizing a biased pre-trained LLM for mining for most gender-biased samples in a dataset, our methods can mitigate gender bias with very few training samples. Finally, we perform an extensive evaluation of our debiasing technique on two recent bias benchmarks (Nadeem et al., 2021) and show that our method outperforms three existing state-of-the-art techniques and performs comparably to the other two. Our main contributions are the following:

- We propose simple data intervention techniques that can be used to reduce gender bias in a pre-trained LLM with few training examples (few-shot), thus making human-in-the-loop bias mitigation strategies feasible.
- We introduce a novel data sampling technique that utilises LLMs to mine for the most biased samples from a dataset and can benefit existing state-of-the-art debiasing methods. When used for debiasing a model, these few samples serve as exemplars and induce large reductions in gender bias.

\* Equal Contribution

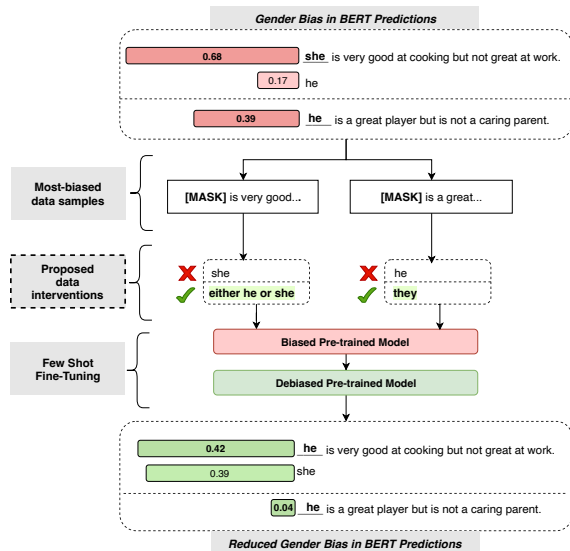


Figure 1: Our method can be summarized as a combination of bias discovery and mitigation. First, we use a pre-trained LLM to find the most gender-biased samples. Then, we apply our data intervention techniques and use these modified training samples to fine-tune the model. Experiments show that our method is very effective at reducing gender bias, outperforming three state-of-the-art baselines and being comparable to two other baselines.

## 2 Related Work

In recent years, there has been growing concern about the bias/stereotypical discriminatory behavior by NLP models, particularly concerning gender. Several studies have investigated the presence of gender bias in various NLP tasks and proposed methods for mitigating it.

One line of research has focused on analyzing the extent of gender bias in pre-trained language models such as BERT and GPT-2. These studies have found that these models exhibit a significant amount of gender bias in their word embeddings for BERT (Jentsch and Turan, 2022) and for GPT-2 (Kirk et al., 2021) and are prone to making stereotypical gender-based predictions (e.g., assuming that a doctor is male and a nurse is female). A standard evaluation metric used in this line of research is Stereotype metrics such as StereoSet (Nadeem et al., 2021), which evaluates the model’s ability to predict gender stereotypes and CrowS pairs (Nangia et al., 2020b) which measure whether a model generally prefers more stereotypical sentences. A similar line of work is gender bias tests proposed in BIG-bench (Srivastava et al., 2022). The tests assess the language model’s gender biases, stereotypes, and ability to infer gender information. It evaluates gender bias and stereotype between male

and female, and gender minority bias and stereotype between majority and minority. It also examines the model’s language modeling performance, which can be affected during de-biasing.

Another line of research has proposed methods for debiasing these models. These methods can be broadly categorized into two groups: **data-based** and **algorithm-based**. Data-based methods aim to reduce bias by removing or altering biased words from the training set. In contrast, algorithm-based methods aim to modify the model’s architecture or training procedure to reduce bias. One popular data-based method is "uncertainty sampling" (Lewis and Gale, 1994), where the model is trained on the instances that it is most uncertain about, which can help to reduce bias by forcing the model to learn from a diverse set of examples. A popular algorithm-based method is "Adversarial Debiasing" proposed by Zhang et al. (2018), which fine-tunes the model using an adversarial loss to make it less sensitive to sensitive attributes such as gender. OSCar proposed by Dev et al. (2021), is another algorithm based method that utilizes the idea of disentangling "problematic concepts" like occupation and gender relationship instead of removing them altogether. MABEL (He et al., 2022) has both algorithm and data-based components, as it first augments the training data by swapping gender words and then applies a contrastive learning objective and alignment via entailment pairs. Their data augmentation strategy is similar in spirit to the data intervention techniques we propose, however our analysis does not require training auxiliary models and uses significantly lesser data.

Data-based methods include the "Equalization" technique proposed by Bolukbasi et al. (2016), which aims to equalize the representation of gender-specific words in the embedding space, the "Counterfactual Data Augmentation" (CDA) method proposed by Zimmermann and Hoffmann (2022), which generates counterfactual examples to improve the model’s robustness to bias, and "Name-Based Counterfactual Data Substitution" proposed by Maudslay et al. (2019) which reduces gender bias by replacing gender-informative names in the dataset with gender-neutral names. Our proposed method is also a data-based method, which aims to effectively reduce gender bias by taking inspiration from different techniques such as uncertainty sampling and name-based counterfactual data substitution (Maudslay et al., 2019).

Gender-Word Pairs	Mean Confidence Difference	
	Mean	Std. Dev.
	he, she	0.317
Will, May	0.316	0.225
boy, girl	0.219	0.218

Table 1: Confidence difference for the Top 3 gender-word pairs in StereoSet

### 3 Probing Bias in Large Language Models

Pre-trained LLMs are biased towards different genders, as seen in a simple mask-fill experiment using BERT. (Here, and in the rest of the paper, we assume a binary treatment of gender for simplicity.) The task is then to mask out the gender-related nouns and pronouns (such as he, she, her, woman, etc.) and get BERT to predict the masked words for the affected sequences in the dataset. Here, we consider a fixed list of gender-specific words curated from previous work (Lu et al., 2018; Zmigrod et al., 2019) and neutral words list<sup>1</sup>. We finally compute the "total confidence difference" as the sum of differences in the model's prediction confidence for each gender-word pair (such as confidence of predicting he – she, man – woman, etc.). Formally, we define total confidence difference as  $|\sum_{i=0}^N (f(x_{female}^{(i)}) - f(x_{male}^{(i)}))|$  where  $f(x)$  represent the confidence of model's prediction,  $N$  is the total number of tokens in the dataset and  $x$  is the tokenized gender word. The higher this number, the more biased the model is concluded to be. We compute the metric at token level and ensure that each of the gender word gets tokenized into exactly one token by initially extending the tokenizer with our gender word list. The top 3 biased gender-word pairs in StereoSet are shown in Table 1. Intuitively, our technique for gauging bias in LLMs is sensitive to the fixed word list used to represent the sensitive attributes (here, gender). In Table 2, we show the number of words covered by the word list used for both WikiText-2 and StereoSet datasets.

### 4 Data Interventions

In order to reduce gender bias in pre-trained models, we carefully select diverse and hard-biased examples and then replace gender words with more neu-

<sup>1</sup><https://github.com/joelparkerhenderson/inclusive-language>

Dataset	Samples	Affected Words (mean)
WikiText-2	10	191
	50	627
	100	1028
StereoSet	10	55
	50	227
	100	463

Table 2: Number of words (mean ) covered by the word list vs dataset and number of sequences sampled from each dataset

tral or equality-focused phrases. This is achieved by using a wordlist to find gender terms in sentences and then segregating words as name and non-name words.

We call our initial approach naive-masking as it does not require a word list for mapping gender words to gender-neutral words. Instead, it replaces all gender words with the fixed word "person." In our next approach, neutral-masking, we swap words in a slightly more semantically accurate manner. In this, we use a word-pair list that goes from gender words to gender-neutral words. With both approaches, we intend to introduce new words in a model's vocabulary to make it more likely to choose a more neutral word in gender-biased sentences.

In our final approach, we exploit the existing vocabulary of the model and try to balance the confidence of prediction on opposite-gender words by using phrases instead. Thus, we call our final approach random-phrase-masking as we instead substitute words with phrases that reflect the equality of gender. This approach not only reduces gender bias but also preserves the original meaning of the sentence in most cases. In our approach, we chose the phrases and order of gender words at random with equal probability.

Intervention	Input word	Converted word
naive-masking	he	person
	she	person
	boy	person
neutral-masking	he	they
	her	their
	schoolgirl	schoolkid
random-phrase-masking	he	he or she
	she	she and he
	boy	either girl or boy

Table 3: Example conversions for three methods. In Random Phrase Masking, the phrase is being chosen and its order was random.

Additionally, we hypothesize that the choice of the dataset for fine-tuning is also essential. We choose two datasets: the WikiText-2 (Merity et al., 2017) dataset, which has implicit gender bias since its sources from Wikipedia articles, and the StereoSet dataset (Nadeem et al., 2021), which has explicit/more gender bias as it has been designed to evaluate gender bias. WikiText-2<sup>2</sup> has 600 train articles and roughly 2M tokens while StereoSet<sup>3</sup> (dev) has 2123 samples out of which we only consider 800 samples which are not unrelated. Naturally, our data intervention method should work better on a dataset with training examples with gender bias while being devoid of meaningful gender associations like "She needs a gynecologist," where the gender of the person is important. By testing our method on both datasets, we can understand the sensitivity of our approach to the quality of training samples used.

## 5 Bias Evaluation Metrics

We focus on evaluating the bias of a model while also measuring its language modeling capability. The ideal model would not just be one with the least bias but also one which does not compromise its language modeling performance. The dual estimation of bias and performance of a model was proposed in the StereoSet benchmark (Nadeem et al., 2021), with the Language Modeling Score (LMS) measuring the percentage of times a meaningful token is predicted for the mask as opposed to a meaningless token, the Stereotype Score (SS) measuring the percentage of times the model predicted a stereotypical word as compared to an anti-stereotypical word, and an idealized CAT score (ICAT) combining the LMS and SS score into a single metric. An ideal model has an ICAT score of 100, while the worst biased model has an ICAT score of 0. We additionally evaluate the CrowS-Pairs benchmark (Nangia et al., 2020a), which captures data with greater diversity in both the stereotypes expressed and the structure of sentences (50 is ideal). However, we note that the Crow-S benchmark is much more limited compared to StereoSet (Nadeem et al., 2021) in terms of both the volume and variety of linguistic phenomenon relating to gender bias it covers.

<sup>2</sup>An English language dataset (Creative Commons Attribution-ShareAlike License).

<sup>3</sup>An English language dataset available at [bias-bench](https://bias-bench.github.io/) (Creative Commons Attribution-ShareAlike 4.0 International Public License)

## 6 Experiments

We compare our proposed interventions with five baselines, 4 of which are state-of-the-art methods and the original pre-trained model. Our first baseline is the application of dropouts to neural networks, **Dropout** proposed by (Webster et al., 2020). Next, we consider an algorithmic de-biasing technique **INLP** technique proposed by (Ravfogel et al., 2020). Then, we consider a sentence embedding de-biasing approach **SentenceDebias** (Liang et al., 2020). Finally, we consider a data-based approach **CDA** (Zmigrod et al., 2019) that is closest to our work. For a fairer comparison, we run the baselines with the same size (100) of the training set as our method. For all of our experiments, we consider the "bert-base-uncased" pre-trained model available from HuggingFace. For fine-tuning our model, we select a varying number of most-biased training samples (10, 50, and 100) from the WikiText-2 and StereoSet (we only use the dev set) datasets, as discussed in section 4. We also compare this to a random selection of data points as an ablation study. On the selected dataset, we apply our interventions and obtain the modified dataset, which is then used to fine-tune our pre-trained model using masked language modeling (MLM) loss. The key point is that we only fine-tune the model on the gender words conditioned on the remaining text, significantly reducing the fine-tuning time. We perform ablations on various types of interventions as discussed in Table 7. The model is trained for 30 epochs, with a learning rate of 0.001 and AdamW optimizer. We ran all of our experiments on NVIDIA Tesla T4 GPU on Google Colab for roughly 48 hours. For all experiments, we report the numbers as the mean and standard deviations (6) of 3 different runs. Our experiment code can be found here.<sup>4</sup>

## 7 Results

Table 4 shows the StereoSet and Crow-S scores for our baselines and our best-performing interventions on the WikiText-2 Dataset. In the StereoSet benchmark, we observe that random-phrase-masking obtains lower SS than all other baselines. On the Crow-S benchmark, random-phrase-masking does better than three of the baselines except SentenceDebias which achieves slightly better scores. While random-phrase-masking results in lower SS scores than neutral-masking, it also obtained

<sup>4</sup>[https://github.com/himansh005/data\\_debias](https://github.com/himansh005/data_debias)

Input Sentence:	[MASK] is very caring and kind but not good at what [MASK] does.
Output of Biased Model:	She is very caring and kind but not good at what she does.
Output of De-biased Model:	He is very caring and kind but not good at what he does.
Input Sentence:	[MASK] is very good at cooking but not great at [MASK] work.
Output of Biased Model:	she is very good at cooking but not great at her work.
Output of De-biased Model:	he is very good at cooking but not great at farm work.
Input Sentence:	Being a [MASK] is not easy since [MASK] will have to stay home and take care of [MASK] child.
Output of Biased Model:	Being a mother is not easy since she will have to stay home and take care of the child.
Output of De-biased Model:	Being a father is not easy since one will have to stay home and take care of the child.

Figure 2: Qualitative analysis of our approach on fill-mask task shows that our intervention techniques are able to modify stereotypical sentences. In this example, we prompted a pre-trained bert-base-uncased model and the same pre-trained model debiased using random-phrase-masking with stereotypical sentences and found that our method is successfully able to reduce biased substitutions.

Type	Method	StereoSet Scores			Crow-S Scores (↓)
		SS (↓)	LMS (↑)	ICAT (↑)	
Baselines	None	60.279	84.172	70.300	57.250
	CDA	60.022	83.466	70.892	56.107
	Dropout	60.529	83.811	70.171	55.977
	SentenceDebias	59.221	<b>84.166</b>	71.308	<b>53.817</b>
	INLP	58.205	83.391	70.966	55.727
Ours	random-phrase-masking (10)	59.442	80.312	70.406	54.580
	random-phrase-masking	<b>58.037</b>	78.676	69.949	54.457
	neutral-masking (10)	60.341	83.956	72.548	55.535
	neutral-masking	60.814	83.587	<b>72.213</b>	56.490

Table 4: StereoSet and Crow-S benchmark results on WikiText-2 dataset. A lower SS and Crow-S score means lesser gender bias while higher ICAT and LMS denote better language modelling ability. The number in the parentheses denotes number of training samples and ones without it use 100.

very low LMS scores. We attribute this performance degradation to the blunt substitution of phrases that our method uses, which might lead to odd-sounding sentences. In the Crow-S benchmarks, we see similar behavior and find that random-phrase-masking does better than neutral-masking. Since we believe that our method is sensitive to the choice of the dataset, we also present results on the StereoSet (dev) dataset 6. In Figure 2, we perform a qualitative analysis of our proposed approach and find that random-phrase-masking is able to flip the predictions on fill-mask tasks for stereotypical sentences.

## 8 Conclusion

In this paper, we show that simple data interventions on limited training data effectively reduce gender bias in LLMs. We also show that a biased pre-trained LLM can be used to mine the most effective de-biasing training examples. Evaluation of our methods on state-of-the-art bias benchmarks empirically suggests that our methods effectively reduce gender bias. Given that our methods can work in a few-shot manner and do not require any auxiliary model training, we hope that our work benefits further research in the domain of human-in-the-loop bias mitigation techniques by making the creation of bias mitigation datasets feasible.

## 9 Limitations

Our proposed method has the following main limitations which we believe are important directions for future work to address:

1. **Gender dependency:** Our approach does not account for sentences that only make sense for a single gender. For example, sentences like "She needs to see a gynecologist" would not be captured by our method. This is a common problem encountered by most debiasing algorithms as it is difficult to distinguish these.
2. **Finite wordlist:** The wordlist does not contain all gender-based words as the language con-

tinues to evolve. We believe that future works could employ better approaches that can automatically mine gender words relevant to a dataset.

3. **Blunt substitution:** The phrase substitution method is an improvement over direct word substitution, but there are still plenty of instances where the new sentence might be semantically incorrect. This does not have any major implication on inference as we are only doing few-shot learning, but it should not be extended to the entire dataset.
4. **Binary gender:** The method only focuses on the male and female gender. It does not consider non-binary or gender-neutral pronouns such as "ze/hir." This can be solved by using an updated wordlist, but the authors could not come across one at the time of writing.
5. **Downstream analyses:** While our work proposes methods that show reduced gender bias as per a set of metrics, the work in no way claims to reduce gender bias in general, especially on downstream tasks. However, we strongly believe that this technique holds potential to reduce gender bias on downstream tasks as well since we adopt a regular fine-tuning approach and focus mainly on better data interventions. Moreover, recent research has shown that fine-tuning-based debiasing approaches do not damage a model's internal representations to a critical extent (Meade et al., 2022).

Overall, these limitations suggest that our approach may not be suitable for use in contexts where gender-specific or non-binary language is prevalent, and the underlying wordlist should be frequently updated.

## 10 Ethics Statement

This study was conducted in accordance with ethical principles and guidelines. The study was designed to provide beneficial knowledge and not harm any group or individual. We recognize that the wordlist we use might not represent all contexts of gender bias and that our debiasing method does not cover all contexts of occurrences of gender bias. However, we made sure to consider the ethical implications of our methodologies and the results of our analysis. The authors have tried to ensure the method

does not amplify any other inherent bias but also acknowledge that our approach may have limitations. We take responsibility for any ethical concerns that may arise as a result of our research.

## Acknowledgments

This material is based upon work partially supported by the National Science Foundation (Awards #1722822 and #1750439) and National Institutes of Health (Awards #R01MH125740, #R01MH096951, and #U01MH116925). PPL is partially supported by a Facebook PhD Fellowship and a Carnegie Mellon University's Center for Machine Learning and Health Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF, NIH, Facebook, or CMLH, and no official endorsement should be inferred. Additionally, we express our appreciation to the anonymous reviewers for their insightful suggestions, which greatly improved our work. Furthermore, we would like to acknowledge the contributions of our colleagues, Atishay Jain and Praneetha Vaddamanu, who played a significant role in the development of this research.

## References

- Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: Allocative versus representational harms in machine learning. In *9th Annual conference of the special interest group for computing, information and society*.
- Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. *Man is to computer programmer as woman is to homemaker? debiasing word embeddings*. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4349–4357.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. *Universal sentence encoder for English*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Sunipa Dev, Tao Li, Jeff M Phillips, and Vivek Srikumar. 2021. *OSCaR: Orthogonal subspace correction and rectification of biases in word embeddings*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5034–

- 5050, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jacqueline He, Mengzhou Xia, Christiane Fellbaum, and Danqi Chen. 2022. [Mabel: Attenuating gender bias using textual entailment data](#). *ArXiv preprint*, abs/2210.14975.
- Sophie Jentzsch and Cigdem Turan. 2022. [Gender bias in BERT - measuring and analysing biases through sentiment rating in a realistic downstream classification task](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 184–199, Seattle, Washington. Association for Computational Linguistics.
- Hannah Rose Kirk, Yennie Jun, Filippo Volpin, Haider Iqbal, Elias Benussi, Frédéric A. Dreyer, Aleksandar Shtedritski, and Yuki M. Asano. 2021. [Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 2611–2624.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. [Measuring bias in contextualized word representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.
- David D. Lewis and William A. Gale. 1994. A sequential algorithm for training text classifiers. In *SIGIR '94*, pages 3–12, London. Springer London.
- Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. [Towards debiasing sentence representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5502–5515, Online. Association for Computational Linguistics.
- Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. [Towards understanding and mitigating social biases in language models](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 6565–6576. PMLR.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2018. [Gender bias in neural natural language processing](#).
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2019. [It’s all in the name: Mitigating gender bias with name-based counterfactual data substitution](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5267–5275, Hong Kong, China. Association for Computational Linguistics.
- Nicholas Meade, Elinor Poole-Dayana, and Siva Reddy. 2022. [An empirical survey of the effectiveness of debiasing techniques for pre-trained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1878–1898, Dublin, Ireland. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pretrained language models](#). In *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)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020a. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020b. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. [Null it out: Guarding protected attributes by iterative nullspace projection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. [Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in NLP](#). *Transactions of the Association for Computational Linguistics*, 9:1408–1424.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2022. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *arXiv preprint arXiv:2206.04615*.
- Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, Ed Chi, and Slav Petrov. 2020. [Measuring and reducing gendered correlations in pre-trained models](#).

Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. [Mitigating unwanted biases with adversarial learning](#). *ArXiv preprint*, abs/1801.07593.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. [Gender bias in contextualized word embeddings](#). 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 629–634, Minneapolis, Minnesota. Association for Computational Linguistics.

Victor Zimmermann and Maja Hoffmann. 2022. [Absinth: A small world approach to word sense induction](#). In *Proceedings of the 18th Conference on Natural Language Processing (KONVENS 2022)*, pages 121–128, Potsdam, Germany. KONVENS 2022 Organizers.

Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach, and Ryan Cotterell. 2019. [Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, Florence, Italy. Association for Computational Linguistics.

## A Appendix

### A.1 Dataset Bias Analysis

To gauge the feasibility of using a wordlist based intervention approach, we first analyze our datasets for occurrences of gender words. As shown in the word cloud 4, gender pronouns are the most frequent word in our datasets. Moreover, as per Figure 1, "she," "he," and "her" are the top three most frequently occurring words in our dataset. This suggests that we can definitely detect gender words in our corpus and apply our interventions.



Figure 3: Frequency of gender words on the StereoSet dataset.




Figure 4: Top 10 most frequent gender words on the StereoSet dataset.

### A.2 Sensitivity to Choice of Dataset

To understand the effectiveness of our proposed data-interventions, we study apply our methods to two datasets under varying number of training samples (10, 50 and 100) and selection strategies (most biased first and random) as per Table 6. Our methods obtain better results on StereoSet (dev) dataset. One reason this could happen is due to the fact that StereoSet has explicit gender bias, thus it would be less likely for a sentence like "She needs a gynaecologist" to appear on it. Because our interventions perform blunt substitutions, this sentence might become incorrect due to our method - "Either he or she needs a gynaecologist".

### A.3 Sensitivity to Number of Training Samples and Sampling Strategy




Figure 5: StereoSet Gender SS scores on StereoSet (dev) dataset with varying number of training samples. Compared to random selection of samples, selecting most-biased samples help achieve lower SS scores.

As per Figure 5, When we vary the number of training samples, we observe that the difference in performance is not huge when we transition from 10 to 100 samples, thus suggesting that our method



Sentences	Mean Confidence Difference
She rushed to see what he wanted and said she loved him. She punched him in the face and told him to go away.	6.85
Jessica is a new mommy. Jessica finds being a mother does not come easy to her. She will no longer work so she can stay home and take care of her child.	6.34
The little girl missed her mommy. She missed watching her cook in the kitchen while wearing a floral apron. She was never home because she worked long hours in the oil field.	4.70

Table 5: Sentences from StereoSet with maximum difference in confidence of prediction between opposite gender words.

Dataset	Sampling Method	Number of Samples	Crow-S Pair Score			StereoSet Scores			Perplexity
			Total	Stereotype Score	Anti-stereotype Score	SS (gender)	LMS (gender)	ICAT (gender)	
StereoSet	most-biased	10	54.481 (2.583)	50.408 (5.295)	60.991 (3.854)	58.736 (1.215)	80.858 (2.988)	66.708 (2.584)	50.449 (54.983)
	random	10	55.47 (3.247)	50.527 (3.632)	63.107 (4.234)	58.952 (0.859)	80.226 (2.85)	65.862 (2.655)	86.024 (107.709)
	most-biased	50	<b>52.994 (1.894)</b>	47.567 (4.564)	61.428 (5.25)	58.498 (1.19)	80.255 (2.428)	66.595 (2.207)	29.599 (28.648)
	random	50	53.817 (1.011)	<b>50.107 (2.972)</b>	59.547 (3.925)	58.485 (0.758)	79.158 (1.992)	65.707 (0.886)	62.498 (11.593)
	most-biased	100	53.054 (2.402)	49.063 (6.025)	59.291 (4.663)	58.071 (1.158)	81.086 (3.226)	67.972 (2.671)	<b>19.079 (14.095)</b>
	random	100	53.563 (1.801)	48.113 (6.499)	62.137 (5.405)	<b>57.719 (1.94)</b>	79.038 (1.406)	66.805 (2.074)	34.826 (12.109)
WikiText-2	most-biased	10	55.6 (3.06)	54.668 (5.606)	57.118 (1.671)	59.344 (0.742)	84.624 (2.134)	<b>68.811 (2.176)</b>	87.06 (80.998)
	random	10	56.617 (1.344)	57.983 (1.305)	<b>54.693 (2.019)</b>	60.616 (0.72)	<b>85.076 (0.896)</b>	67.021 (1.895)	59.901 (102.019)
	most-biased	50	54.276 (1.513)	53.394 (3.847)	55.834 (2.652)	59.238 (1.068)	83.348 (3.003)	67.902 (0.977)	212.365 (155.526)
	random	50	54.2 (2.383)	51.783 (5.272)	57.93 (2.969)	59.611 (1.155)	83.456 (1.9)	67.386 (0.74)	116.872 (100.401)
	most-biased	100	55.473 (1.42)	54.827 (4.255)	56.637 (4.329)	59.426 (1.719)	83.442 (3.185)	67.629 (1.178)	220.957 (207.243)
	random	100	54.457 (1.444)	51.363 (4.283)	59.223 (4.451)	59.545 (0.387)	81.953 (1.442)	66.3 (0.597)	326.017 (181.822)

Table 6: StereoSet and Crow-S scores for random-phrase-masking method on two datasets, 3 sample sizes and 2 selection methods. We report mean (standard deviation) across 3 different runs. Selecting most-biased samples and using the StereoSet dataset for fine-tuning gives best results.

Name Word Mask Method	Non-Name Word Mask Method	Crow-S Pairs			StereoSet			Perplexity
		Total	Stereotype Score	Anti-Stereotype Score	SS	LMS	ICAT	
	female-first-random-phrase-masking	53.18 (3.106)	49.06 (6.627)	59.55 (3.678)	58.283 (0.4)	79.059 (0.436)	65.96 (0.27)	26.3 (9.545)
	naive-masking	<b>50.637 (0.585)</b>	43.607 (1.449)	61.49 (1.486)	59.521 (0.458)	83.325 (0.62)	67.456 (0.414)	<b>1.0 (0.0)</b>
naive-masking	random-phrase-masking	52.673 (1.374)	49.057 (5.998)	58.253 (5.906)	58.05 (0.851)	78.218 (0.633)	65.618 (0.937)	30.045 (8.019)
neutral-masking	female-first-random-phrase-masking	53.44 (0.0)	53.46 (0.891)	<b>53.4 (1.372)</b>	58.246 (0.285)	<b>87.182 (0.391)</b>	<b>72.806 (0.823)</b>	11.39 (6.649)
	random-phrase-masking	54.195 (1.619)	48.43 (0.891)	63.11 (2.744)	57.316 (0.164)	78.339 (0.196)	66.877 (0.424)	54.413 (0.212)
	fixed-phrase-masking-1	53.307 (1.761)	46.837 (8.494)	63.43 (8.807)	57.688 (1.718)	79.554 (0.17)	67.32 (2.64)	14.484 (1.512)
	fixed-phrase-masking-2	51.783 (4.965)	46.43 (10.381)	60.193 (3.503)	57.229 (1.739)	80.551 (1.251)	68.879 (1.882)	13.374 (1.174)
	fixed-phrase-masking-3	52.927 (1.541)	48.317 (3.78)	60.193 (4.234)	<b>56.963 (1.373)</b>	79.3 (1.531)	68.284 (3.478)	15.546 (2.997)
	fixed-phrase-masking-4	53.567 (4.186)	<b>50.083 (9.006)</b>	59.223 (3.885)	58.13 (1.208)	79.834 (0.533)	66.86 (2.309)	14.51 (1.339)

Table 7: StereoSet Gender SS scores on StereoSet (dev) dataset on 100 samples across various interventions techniques. All numbers are reported as mean and standard deviation across 3 runs.

is capable of few-shot fine-tuning. Moreover, sampling the most biased data points helps our methods achieve better performance consistently, as shown in Figure 5 and Table 6. Table ?? shows some top three most gender biased entries found in the StereoSet dataset.

#### A.4 Ablations of interventions

We study the effects of choosing different ways of replacement for name and non-name words. In addition to our three interventions proposed previously, we also experimented with a couple of others. In `female-first-random-phrase-masking`, we always keep the female gendered word before a male word. We wanted to understand if the order of gender words encountered by a model renders any effect on the debiasing. In Table 7, we see that it does not perform any better than `random-phrase-masking`. Then, we also try fixing the phrases from `random-phrase-masking`, thus making it `fixed-phrase-masking`. We obtain 4 variants of this method corresponding to the following four phrases:

1. both [1] and [2]
2. [1] and [2]
3. [1] or [2]
4. either [1] or [2]

Here, [1] and [2] are substituted with opposite gender words. As we observe in Table 7, `fixed-phrase-masking-3` obtains the lowest StereoSet Gender SS score out of all our intervention methods. Similarly, `naive-masking` obtains the lowest Crow-S pair score.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

A1. Did you describe the limitations of your work?

9

A2. Did you discuss any potential risks of your work?

10

A3. Do the abstract and introduction summarize the paper's main claims?

1

A4. Have you used AI writing assistants when working on this paper?

*Left blank.*

### B Did you use or create scientific artifacts?

4

B1. Did you cite the creators of artifacts you used?

4

B2. Did you discuss the license or terms for use and / or distribution of any artifacts?

*The datasets used are open-source.*

B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

4

B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?

*The data is a general vocabulary. Hence, none of the used data contains any personal information.*

B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

*The dataset pages have all the needed information*

B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.

7

### C Did you run computational experiments?

6

C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

6

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

6

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

6

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*Not applicable. Left blank.*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*Not applicable. Left blank.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*Not applicable. Left blank.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*Not applicable. Left blank.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*Not applicable. Left blank.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*Not applicable. Left blank.*