

FineDeb: A Debiasing Framework for Language Models

Akash Saravanan*, Dhruv Mullick*, Habibur Rahman*, Nidhi Hegde*

{*asaravan, mullick, habibur, nidhi.hegde*}@ualberta.ca

Department of Computing Science, Alberta Machine Intelligence Institute (Amii)
University of Alberta, Canada

Abstract

As language models are increasingly included in human-facing machine learning tools, bias against demographic subgroups has gained attention. We propose FineDeb, a two-phase debiasing framework for language models that starts with contextual debiasing of embeddings learned by pre-trained language models. The model is then fine-tuned on a language modeling objective. Our results show that FineDeb offers stronger debiasing in comparison to other methods which often result in models as biased as the original language model. Our framework is generalizable for demographics with multiple classes, and we demonstrate its effectiveness through extensive experiments and comparisons with state of the art techniques. We release our code and data on GitHub¹.

1 Introduction & Related Work

Machine learning tools that rely on natural language processing (NLP) are increasingly being developed for scenarios with immediate impact on individuals, such as health-care (Velupillai et al. 2018), conversational agents (Zhang et al. 2020), and legal systems (Dale 2019). While the language models here vary in the type of embeddings used, they rely on representations that may reflect or exhibit bias (Manzini et al. 2019; Bolukbasi et al. 2016). When used in downstream tasks such as prediction, health care diagnoses, or other decision-making, such representations can amplify bias and result in discriminatory actions against individuals in disadvantaged demographic subgroups. Our focus in the present paper is on such representational harms (Blodgett et al. 2020).

Prior work studies word embeddings to analyze bias and proposes debiasing techniques for NLP methods. Debiasing on word embeddings was first introduced by Bolukbasi et al. (2016) and further refined to enable debiasing on multiple classes by Manzini et al. (2019). However, recent advances in NLP have been focused on large pretrained transformer-based language models (LM) such as BERT and GPT. Such models differ in that rather than considering individual word embeddings, they create representations that take into account large and connected components such as sentences

and context. For comprehensive bias mitigation we must consider bias in the context of sentences, beyond mere word embeddings. We therefore focus on bias in transformer-based contextual LMs.

Previous work (Bolukbasi et al. 2016; Liang et al. 2020, 2021) applies debiasing techniques on embeddings after bias subspaces in these representations have been detected. While such techniques attempt to create debiased representations, downstream tasks may not necessarily reflect debiased language. Furthermore, good performance on those downstream tasks, such as entity extraction or text classification, must still be an objective. Bordia and Bowman (2019) thus act on the training objective by adding a regularizer for debiasing. In place of a regularizer, we use an entirely new training objective to minimize distance between relevant embeddings.

A further limitation of state of the art debiasing techniques is that they largely consider demographics with binary classes or when they consider multiple classes, they focus on the three largest subgroups (Manzini et al. 2019). In real scenarios, social disadvantage is represented through more than the binary or majority/minority dichotomy, with demographic groups containing many classes. We apply our debiasing methodology on demographics with multiple classes.

We propose a debiasing framework - FineDeb for large pretrained LMs, with two phases: one for debiasing the embeddings, and one for fine-tuning the language modeling head. **1)** In the debiasing phase we modify representations by using a training objective that minimizes the distance between embeddings of target words while considering their sentence contexts. **2)** In the second phase, we fine-tune the model on a language modeling task.

Our results show that the FineDeb framework applies a significantly stronger level of debiasing in comparison to other methods while slightly compromising on downstream language modeling performance. In this paper we extend prior work in several ways: compiling lists of target words for three demographics, generating a debiasing dataset for future benchmarking, providing a framework for debiasing language models rather than directly modifying representations, demonstrating a debiasing training objective, and allowing for the debiasing of multiple classes.

*These authors contributed equally.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/akashsara/debiasing-language-models>

2 Data

We use two types of data - one for the debiasing phase and one for the LM finetuning phase.

2.1 Debiasing Data

Our methodology starts by first finetuning the LM with a new objective function on a debiasing dataset, which consists of examples of debiased sentences. The debiased sentence examples are created using word lists crafted with multiple classes per demographic (in english language). Our final debiased model is thus trained by debiasing for multiple classes. Note that this is a novel contribution beyond word lists with pairs created previously.

Our word lists are compiled from various sources, both online (A) and existing work (Bolukbasi et al. 2016; Zhao et al. 2018), to create a list of target word tuples for each demographic (race, religion and gender). The word list contains 2 classes for gender², 5 classes for race, and 7 classes for religion. For each demographic, the word list consists of several tuples of target words, where within each tuple the words are comparable. For example, ("*Muslim*", "*Christian*", "*Jewish*", "*Hindu*", "*Buddhist*", "*Confucianist*", "*Taoist*") is one tuple within the religion word list. We compile 10 word tuples for Race, 32 for Religion and 158 for Gender. We make the word lists freely available in our codebase, providing samples in Appendix A.

In order to generate our final debiasing dataset, we first craft sentence templates from the RedditBias (Barikeri et al. 2021) dataset. RedditBias is a dataset of human conversation data from Reddit across four demographics: *gender*, *race*, *religion*, and *queerness*. We convert all sentences containing the former three demographics into sentence templates by removing the target words. An example for race would be "*all ____ are criminals*". During the training process we choose a relevant word tuple and sample target words from different classes such as *Black* and *White* to generate sentence pairs that differ in only the target word. We generate such sentence sets pairwise among all classes.

2.2 LM finetuning Data

Taking inspiration from prior work (Qian et al. 2019; Bordia and Bowman 2019), we use CNN-DailyMail (Hermann et al. 2015) for our language model finetuning objective. It consists of 300,000+ English news articles.

3 FineDeb

Our method, FineDeb, adopts a two phase approach for training. In the first phase, we debias the model by modifying the embeddings learned by the language model, and in the second phase, we finetune the debiased model on the language modeling objective. We evaluate our framework in terms of bias and downstream language modeling performance using standard metrics. Our method is demonstrated on a BERT model (Devlin et al. 2019), specifically *bert-base-uncased*. All hyperparameters are listed in Appendix D.

²NLP datasets have limited data for other gender classes

3.1 Debiasing Phase

In the debiasing phase, we train our model using the sentence pairs generated in Section 2.1. Most modern language models are contextual in nature (Devlin et al. 2019), that is, the same word may have different meanings based on the context in which the word is used. For instance, the word "*temple*" could refer to a building or a part of the human body. Thus it becomes important to perform our debiasing only in those contexts where the bias may exist. Our methodology is inspired by the traditional process of determining relationships between pairs of words by computing the distance between their embeddings (Mikolov et al. 2013; Bolukbasi et al. 2016). However, unlike these methods, we account for the contextual meaning of the sentence.

Given two near-identical sentences that differ only by a target word (or phrase), we first compute the difference between the embeddings of the two sentences (using their [CLS] token embedding as in Devlin et al. (2019)), and the difference between the embeddings of the target words in each sentence. Our training objective is motivated by the idea that since the two sentences differ only in their target words, logically, the difference in the embeddings of these two sentences should also differ only by the difference in the embeddings of the target words. Our training objective thus minimizes the distance between these two quantities to debias the model. Formally, our loss function is as follows: $L(S_1, S_2, W_1, W_2) = \mathcal{D}(S_1 - S_2, W_1 - W_2)$, where S_i is the embedding for sentence i , W_i is the embedding for the target word in sentence i , and $\mathcal{D}(\cdot)$ is the distance between the two quantities (we use the Mean Squared Error but other similarity measures such as Cosine similarity can also be used). Taking an example from Nangia et al. (2020), if we had two sentences "*The crafty Jews made a plan to steal the money.*" and "*The crafty Christians made a plan to steal the money.*", W_1 and W_2 would be the embeddings of "*Jews*" and "*Christians*" respectively, while S_1 and S_2 would respectively be the sentence embeddings.

3.2 LM Finetuning Phase

Pretrained transformer models consist of an embedding-generating model, and a language modeling head (LM head) which gives probabilities for each of the words based on these embeddings. In the case of our debiasing phase, we only update the embedding-generating model and not the LM head. We then finetune the entire model (debiased model + LM head) on the standard BERT language modeling objective (masked word prediction) using the CNN-DailyMail dataset.

During this finetuning process, we freeze the debiased model so that the debiased embeddings do not change, and only the weights in the LM head are updated. This results in an overall model suitable for the LM task. While this method of improving the language modeling ability of our model may re-introduce some biases that exist in the CNN-DailyMail dataset, it does not eliminate the effects of our debiasing, as we show in our results. Indeed the two phases may counteract each other, but this tradeoff can be controlled via the amount of training done in each phase.

4 Metrics

We evaluate FineDeb on the three demographics of gender, race, and religion using three metrics: StereoSet, SEAT, and Crow-S Pairs. We choose these metrics due to their popularity in prior work (Meade, Poole-Dayana, and Reddy 2021), though we note that other metrics may exist (Kurita et al. 2019). These metrics differ in how bias is evaluated, the data used in evaluating bias, and whether the language modeling performance is considered.

4.1 StereoSet

Following recent work (Meade, Poole-Dayana, and Reddy 2021; Zhang et al. 2022), we use StereoSet (Nadeem, Bethke, and Reddy 2021) to evaluate our work. StereoSet measures the Stereotype Score (SS) which gives a measure of bias and the Language Modeling Score (LMS) which determines performance at language modeling tasks. There is a trade-off here, as a model could be perfectly unbiased but be a poor language model (or vice versa). Thus the authors provide a combined measure named ICAT. Following prior work (Meade, Poole-Dayana, and Reddy 2021), we use the intrasentence variant of StereoSet.

4.2 Crow-S Pairs

The Crowdsourced Stereotype Pairs (Crow-S Pairs) (Nangia et al. 2020) uses crowdsourced pairs of sentences that differ only by a small number of tokens such that one sentence reflects a stereotype while the other violates that stereotype. Under this metric, a perfect model is equally likely to pick the stereotypical sentence as it is to pick the anti-stereotypical sentence. This metric does not test the language modeling ability of the model but covers a wide variety of biases.

4.3 SEAT

Sentence Encoder Association Test (SEAT) (May et al. 2019) is a sentence level extension of the WEAT metric (Caliskan, Bryson, and Narayanan 2017) which measures bias between two sets of attribute words and two sets of target words. Specifically, SEAT uses sentence templates to obtain representations of words. The metric is measured in terms of the average effect size across several tests, where a value closer to 0 indicates a lower degree of bias but it does not test the language modeling ability of the model.

5 Results and Discussion

We compare FineDeb on three demographics against two baselines and five prior works. Our baselines are pretrained BERT and a pretrained BERT model where only the LM finetuning phase is applied. We include results on this latter model to show that the LM finetuning phase does not significantly alter bias in the model compared to the base BERT model, and any change in bias in our model is strictly due to our debiasing phase. This is evident in the results which show that BERT with LM finetuning is on par with or slightly better than the base BERT model for all listed metrics. We also compare with state of the art methods - CDA,

Dropout, INLP, Self-Debias and, Sentence Debias, the results of which we cite from Meade, Poole-Dayana, and Reddy (2021).

We first present results³ on the StereoSet metric in Table 1. Under the SS measure, FineDeb outperforms all techniques for all three demographics. Furthermore, the other debiasing methods have small improvements in comparison to the baseline BERT model while our work shows near-perfect scores for two of the three demographics (race and religion). Under the LMS measure, Self-Debias, Sentence Debias, and the baselines have the best or second best results, across different demographics. It is expected that our method not perform best on the LMS measure since we focus first and foremost on debiasing. However the reduction in LMS performance is not too severe since under the combined measure of ICAT our model performs best for gender and religion, and second best for race.

For SEAT (Table 2), our method performs second best for gender, while for race and religion we do better than one method and no other methods respectively. We note that SEAT measures distances between embeddings, whereas StereoSet is based on final word outputs. Further, SEAT does not measure LM performance. We reason that since disadvantage and discriminatory harm is manifested due to contextual outputs from language models and not just their representations, a metric such as StereoSet, which considers both a measure of final LM performance, as well as a measure of the bias in contextual outputs, more accurately represents real-world impact of bias.

The results for Crow-S (Table 3) show that, FineDeb outperforms all others, with similar performance as Dropout for religion. For gender, our method beats the baseline and two other methods, but performs poorly for race. Crow-S measures bias similarly to StereoSet by considering whether a stereotypical sentence is preferred among sentence pairs. The sentences in a pair, however, differ on attributes rather than target words. We cannot be certain if this difference accounts for the differing results, but we note that like SEAT, Crow-S does not measure LM ability, an important component of any model.

6 Conclusions and Future Work

In this paper, we have proposed FineDeb, a new framework for debiasing language models and have demonstrated its performance on commonly used evaluation metrics. Considering the results over three demographics and across three metrics against all benchmarks, we see that our FineDeb framework is strongly debiasing. This could be since the training objective itself minimizes embedding distance. At the same time, this may lead to a distortion in the embeddings. This is seen in the near-perfect StereoSet SS measures and in the somewhat poorer scores for SEAT which directly works with the embeddings. Since the other debiasing solutions have an SS score similar to the default BERT case, their debiasing ability is not as strong as FineDeb. For the

³Evaluation code taken from Nadeem, Bethke, and Reddy (2021) (StereoSet) and Meade, Poole-Dayana, and Reddy (2021) (SEAT and Crow-S)

| StereoSet | Gender | | | Race | | | Religion | | |
|----------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | LMS | SS | ICAT | LMS | SS | ICAT | LMS | SS | ICAT |
| BERT | 84.17 | 60.28 | 66.86 | <u>84.17</u> | 57.03 | 72.34 | 84.17 | 59.70 | 67.84 |
| BERT (LM finetuning) | 85.01 | 59.01 | <u>69.69</u> | 83.78 | 56.38 | 73.08 | 83.07 | 61.47 | 64.01 |
| FineDeb | 77.70 | 53.27 | 72.62 | 75.37 | 50.82 | <u>74.13</u> | 74.10 | 50.39 | 73.52 |
| CDA | 83.08 | 59.61 | 67.11 | 83.41 | 56.73 | 72.18 | 83.24 | 58.37 | 69.31 |
| Dropout | 83.04 | 60.66 | 65.34 | 83.04 | 57.07 | 71.30 | 83.04 | 59.13 | 67.88 |
| INLP | 80.63 | <u>57.25</u> | 68.94 | 83.12 | 57.29 | 71.00 | 83.36 | 60.31 | 66.17 |
| Self-Debias | 84.09 | 59.34 | 68.38 | 84.24 | <u>54.30</u> | 77.00 | <u>84.23</u> | <u>57.26</u> | <u>72.00</u> |
| Sentence Debias | <u>84.20</u> | 59.37 | 68.42 | 83.95 | 57.78 | 70.89 | 84.26 | 58.73 | 69.55 |

Table 1: StereoSet evaluation. LMS indicates Language Modeling Score, SS is the Stereotype Score, ICAT is the overall score. Higher is better for LMS & ICAT, closer to 50 is better for SS (**Best**; Next Best).

| | Gender | Race | Religion |
|-----------------|-------------|-------------|-------------|
| BERT | 0.62 | 0.62 | 0.49 |
| FineDeb | <u>0.36</u> | 0.62 | 0.67 |
| CDA | 0.72 | <u>0.57</u> | 0.34 |
| Dropout | 0.77 | 0.55 | <u>0.38</u> |
| INLP | 0.20 | 0.64 | 0.46 |
| Sentence Debias | 0.43 | 0.61 | 0.44 |

Table 2: SEAT: Average effect size for each demographic. BERT (LM finetuning) & Self-Debias do not modify internal representations and so have the same SEAT score as BERT. Lower is better (**Best**; Next Best).

| | Gender | Race | Religion |
|----------------------|--------------|--------------|--------------|
| BERT | 57.25 | 62.33 | 62.86 |
| BERT (LM finetuning) | 57.63 | 62.91 | 58.10 |
| FineDeb | 54.58 | 65.24 | 44.76 |
| CDA | 56.11 | 56.70 | 60.00 |
| Dropout | 55.34 | <u>59.03</u> | 55.24 |
| INLP | 51.15 | 67.96 | 60.95 |
| Self-Debias | <u>52.29</u> | 56.70 | <u>56.19</u> |
| Sentence Debias | <u>52.29</u> | 62.72 | 63.81 |

Table 3: Crow-S Pairs: Metric scores for each demographic. Closer to 50 is better (**Best**; Next Best).

ICAT measure which combines debiasing and LM performance our method performs the best or second best, suggesting a better overall performance for our method.

The tradeoff between the debiasing performance (measured by SS) and language modeling performance (measured by LMS score) which we note in Section 3.2 can be controlled by hyperparameter tuning during training. A deteriorated language modeling ability would have an impact on other downstream NLP tasks (Wang et al. 2019). Depending on the use case, either a more debiased model or a model with better language modeling ability (and consequently better downstream NLP task ability) can be obtained.

We note that despite FineDeb achieving near-perfect debiasing performance at the cost of language modeling performance, the language modeling performance may not always be of primary importance. We envision a number of

scenarios in which the extent of debiasing would be more important than LM performance. It is easier for a human to identify and check for correctness than it is to identify and correct biases, some of which they may not even be aware of. Thus, any human-in-the-loop task would benefit from a model generating suggestions free from biases. For instance, consider the case of generating text descriptions for products Koto, Lau, and Baldwin (2022). A debiased model can first generate an unbiased description, and then a human can fix any errors, if necessary. Similarly FineDeb can play an important role in other human-in-the-loop scenarios like generating stories, or in recommendation systems where a slight decrease in recommendation quality is preferable to serving biased recommendations.

There are a few avenues for future work in this area. Namely, expanding our work to include more classes (such as in gender) or to other demographics; a more comprehensive analysis of our model on downstream tasks and modifying our current two-phase framework approach to instead use a single phase of interleaved debiasing and finetuning. Further, current metrics for bias rely on comparisons of either target words or attribute words, resulting in varying performance across the different techniques. This suggests the need for a more comprehensive metric on bias that is agnostic to the debiasing technique.

7 Limitations

To encourage transparency in research, we disclose some limitations of our work in the hope that it offers further insights our the proposed framework:

- There may be cases where we want a bias to exist. For example, in the sentence "The ___ man went to the mosque.", the probability of "Muslim" should be higher than the probability of "Christian" or "Jew". While people of any religion could go to a mosque, a person who follows Islam is far more likely. This falls under explainable bias (Mehrabi et al. 2021).
- While our method is effective, it relies heavily on the word lists that we have compiled to the best of our knowledge. These are by no means fully representative of all bias targets for a given demographic and there is still scope for expansion.

Appendix

| Male | Female |
|------------|--------------|
| countryman | countrywoman |
| fraternal | sororal |
| manservant | maidservant |
| divo | diva |
| actor | actress |
| bachelor | spinster |
| papa | mama |
| busboy | busgirl |

Table 4: Example word lists for Gender

A Word Lists

We compile our word lists from several online sources in addition to prior work (Bolukbasi et al. 2016; Zhao et al. 2018). Specifically, we consider the following pages on Wikipedia: Major Religious Groups, Race (Human Categorization). In addition, we also consulted <https://courses.lumenlearning.com/sociology/chapter/world-religions/>.

A.1 Gender

In Table 4, we provide a sample of word lists in the Gender word list.

A.2 Race

In Table 5, we provide a sample of word lists in the Race word list.

A.3 Religion

In Table 6, we provide a sample of word lists in the Religion word list.

B Additional Results

We provide the complete evaluation results for SEAT in Tables 7, 8 and 9.

C Training Details

For each model, we run a job with a time limit of 41 hours, 4 CPUs per job, 4GB of memory per CPU and 1 NVIDIA V100 GPU.

D Hyperparameters

Considering the trade-off between bias reduction and LM performance discussed in this paper, we experiment with different dataset sizes for both the debiasing phase and LM finetuning phase. Intuitively, a larger debiasing dataset size would lead to a higher bias reduction but poorer LM performance. On the other hand, a large LM finetuning dataset would lead to higher bias but also better LM performance. The debiasing dataset sampling size is varied as - 500, 1000, 2000, 4000, and continues upto the debiasing dataset’s size (different for different demographics). The LM finetuning

dataset size is varied as a percentage of the total CNN dataset - 1%, 2%, 4%, 8%, 16% and 32%.

For training our language model (3.1, 3.2) we use a batch size of 64 to maximise GPU usage, maximum source length of 64. For language model finetuning, we use 100 epochs to allow for convergence, whereas for debiasing we find 30 epochs to be sufficient, with both training processes utilizing an early stopping mechanism. We use a learning rate of $1e-4$ based on what we find in the original BERT paper (Devlin et al. 2019).

References

- Barikeri, S.; Lauscher, A.; Vulić, I.; and Glavaš, G. 2021. RedditBias: A Real-World Resource for Bias Evaluation and Debiasing of Conversational 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)*, 1941–1955. Online: Association for Computational Linguistics.
- Blodgett, S. L.; Barocas, S.; Daumé III, H.; and Wallach, H. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *ACL*.
- Bolukbasi, T.; Chang, K.-W.; Zou, J.; Saligrama, V.; and Kalai, A. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, 4356–4364. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510838819.
- Bordia, S.; and Bowman, S. R. 2019. Identifying and Reducing Gender Bias in Word-Level Language Models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, 7–15. Minneapolis, Minnesota: Association for Computational Linguistics.
- Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334): 183–186.
- Dale, R. 2019. Law and Word Order: NLP in Legal Tech. *Nat. Lang. Eng.*, 25(1): 211–217.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Hermann, K. M.; Kociský, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching Machines to Read and Comprehend. In *NIPS*, 1693–1701.
- Koto, F.; Lau, J. H.; and Baldwin, T. 2022. Can Pretrained Language Models Generate Persuasive, Faithful, and Informative Ad Text for Product Descriptions? In *Proceedings of the Fifth Workshop on e-Commerce and NLP (ECNLP 5)*,

| | | | | |
|------------------|----------------|--------------|-----------|-----------------|
| African-American | Anglo-American | Hispanic | Asian | Native-American |
| Black | White | Latino | Brown | light-brown |
| Negroid | Caucasian | Latino | Brown | Native-American |
| African-American | Anglo-American | Hispanic | Asian | Native-American |
| Afro-American | Anglo-American | Hispanic | Asian | Native-American |
| African | American | Hispanic | Asian | Native-American |
| Afroamerican | Angloamerican | Hispanic | Asian | Native-American |
| Negro | Caucasian | Hispanic | Brown | Native-American |
| dark-skin | light-skin | white-latino | gray-skin | reddish-brown |
| dark-skin | light-skin | black-latino | gray-skin | reddish-brown |

Table 5: Example word lists for Race

| | | | | | | |
|----------|--------------|---------|-------------|-----------|--------------|--------------|
| Islam | Christianity | Judaism | Hinduism | Buddhism | Confucianism | Taoism |
| Islam | Christianity | Judaism | Hinduism | Buddhism | Ruism | Daoism |
| Quran | Bible | Torah | Gita | Tripitaka | Analects | Tao-Te-Ching |
| Koran | Bible | Tanakh | Veda | Tripitaka | Analects | Tao-Te-Ching |
| Muslim | Christian | Jewish | Hindu | Buddhist | Confucianist | Taoist |
| islamic | Christian | Jewish | Hinduism | Buddhist | Confucius | Dao |
| Mohammed | Jesus | Malachi | Ramakrishna | Gautama | Confucius | Laozi |
| Mohammed | Jesus | Moses | Krishna | Gautama | Kung-Fu-Tzu | Laozi |

Table 6: Example word lists for Religion

| Model | SEAT-6 | SEAT-6b | SEAT-7 | SEAT-7b | SEAT-8 | SEAT-8b | Avg. Effect Size |
|-------------------|--------------|---------|--------|--------------|--------------|--------------|------------------|
| bert-base-uncased | <u>0.931</u> | 0.090 | -0.124 | <u>0.937</u> | <u>0.783</u> | <u>0.858</u> | 0.620 |
| FineDeb | 0.248 | 0.113 | 0.179 | 0.298 | 0.280 | <u>1.056</u> | 0.362 |
| CDA | <u>.846</u> | 0.186 | -0.278 | <u>1.342</u> | <u>0.831</u> | <u>0.849</u> | 0.722 |
| Dropout | <u>1.136</u> | 0.317 | 0.138 | <u>1.179</u> | <u>0.879</u> | <u>0.939</u> | 0.765 |
| INLP | 0.317 | -0.354 | -0.258 | 0.105 | 0.187 | -0.004 | 0.204 |
| Sentence Debias | 0.350 | -0.298 | -0.626 | <u>0.458</u> | 0.413 | <u>0.462</u> | 0.434 |

Table 7: SEAT effect sizes (closer to 0 is better) for Gender debiased models. Underlined if statistically significant ($p < 0.01$).

| Model | ABW-1 | ABW-2 | SEAT-7 | SEAT-3 | SEAT-4 | SEAT-5 | SEAT-5b | Avg. Effect Size |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|------------------|
| bert-base-uncased | -0.079 | <u>0.690</u> | <u>0.778</u> | <u>0.469</u> | <u>0.901</u> | <u>0.887</u> | <u>0.539</u> | 0.620 |
| FineDeb | <u>1.111</u> | -0.235 | <u>0.806</u> | 0.085 | <u>0.858</u> | <u>0.787</u> | <u>0.456</u> | 0.620 |
| CDA | 0.231 | <u>0.619</u> | <u>0.824</u> | 0.510 | <u>0.896</u> | <u>0.418</u> | <u>0.486</u> | 0.569 |
| Dropout | <u>0.415</u> | <u>0.690</u> | <u>0.698</u> | <u>0.476</u> | <u>0.683</u> | <u>0.417</u> | <u>0.495</u> | 0.554 |
| INLP | 0.295 | <u>0.565</u> | <u>0.799</u> | <u>0.370</u> | <u>0.976</u> | <u>1.039</u> | <u>0.432</u> | 0.639 |
| Sentence Debias | -0.067 | <u>0.684</u> | <u>0.776</u> | <u>0.451</u> | <u>0.902</u> | <u>0.891</u> | <u>0.513</u> | 0.612 |

Table 8: SEAT effect sizes (closer to 0 is better) for Race debiased models. Underlined if statistically significant ($p < 0.01$).

| Model | Religion-1 | Religion-1b | Religion-2 | Religion-2b | Avg. Effect Size |
|-------------------|--------------|--------------|--------------|--------------|------------------|
| bert-base-uncased | <u>0.744</u> | -0.067 | <u>1.009</u> | -0.147 | 0.492 |
| FineDeb | <u>0.697</u> | <u>0.701</u> | <u>0.666</u> | <u>0.613</u> | 0.670 |
| CDA | <u>0.355</u> | -0.104 | <u>0.424</u> | -0.474 | 0.339 |
| Dropout | <u>0.535</u> | 0.109 | <u>0.436</u> | -0.428 | 0.377 |
| INLP | <u>0.473</u> | -0.301 | <u>0.787</u> | -0.280 | 0.460 |
| Sentence Debias | <u>0.728</u> | 0.003 | <u>0.985</u> | 0.038 | 0.439 |

Table 9: SEAT effect sizes (closer to 0 is better) for Religion debiased models. Underlined if statistically significant ($p < 0.01$).

- Representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, 166–172. Florence, Italy: Association for Computational Linguistics.
- Liang, P. P.; Li, I. M.; Zheng, E.; Lim, Y. C.; Salakhutdinov, R.; and Morency, L.-P. 2020. Towards Debiasing Sentence Representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5502–5515. Online: Association for Computational Linguistics.
- Liang, P. P.; Wu, C.; Morency, L.-P.; and Salakhutdinov, R. 2021. Towards Understanding and Mitigating Social Biases in Language Models. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 6565–6576. PMLR.
- Manzini, T.; Yao Chong, L.; Black, A. W.; and Tsvetkov, Y. 2019. Black is to Criminal as Caucasian is to Police: Detecting and Removing Multiclass Bias in 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)*, 615–621. Minneapolis, Minnesota: Association for Computational Linguistics.
- May, C.; Wang, A.; Bordia, S.; Bowman, S. R.; and Rudinger, R. 2019. On Measuring Social Biases in Sentence Encoders. 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)*, 622–628. Minneapolis, Minnesota: Association for Computational Linguistics.
- Meade, N.; Poole-Dayana, E.; and Reddy, S. 2021. An Empirical Survey of the Effectiveness of Debiasing Techniques for Pre-trained Language Models.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.*, 54(6).
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. In Bengio, Y.; and LeCun, Y., eds., *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Nadeem, M.; Bethke, A.; and Reddy, S. 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)*, 5356–5371. Online: Association for Computational Linguistics.
- Nangia, N.; Vania, C.; Bhalerao, R.; and Bowman, S. R. 2020. 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)*, 1953–1967. Online: Association for Computational Linguistics.
- Qian, Y.; Muaz, U.; Zhang, B.; and Hyun, J. W. 2019. Reducing Gender Bias in Word-Level Language Models with a Gender-Equalizing Loss Function. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, 223–228. Florence, Italy: Association for Computational Linguistics.
- Velupillai, S.; Suominen, H.; Liakata, M.; Roberts, A.; Shah, A. D.; Morley, K.; Osborn, D.; Hayes, J.; Stewart, R.; Downs, J.; Chapman, W.; and Dutta, R. 2018. Using clinical Natural Language Processing for health outcomes research: Overview and actionable suggestions for future advances. *Journal of Biomedical Informatics*, 88: 11–19.
- Wang, A.; Hula, J.; Xia, P.; Pappagari, R.; McCoy, R. T.; Patel, R.; Kim, N.; Tenney, I.; Huang, Y.; Yu, K.; Jin, S.; Chen, B.; Van Durme, B.; Grave, E.; Pavlick, E.; and Bowman, S. R. 2019. Can You Tell Me How to Get Past Sesame Street? Sentence-Level Pretraining Beyond Language Modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4465–4476. Florence, Italy: Association for Computational Linguistics.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; Mihaylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.; Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models.
- Zhang, Y.; Sun, S.; Galley, M.; Chen, Y.-C.; Brockett, C.; Gao, X.; Gao, J.; Liu, J.; and Dolan, B. 2020. DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 270–278. Online: Association for Computational Linguistics.
- Zhao, J.; Zhou, Y.; Li, Z.; Wang, W.; and Chang, K.-W. 2018. Learning Gender-Neutral Word Embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4847–4853. Brussels, Belgium: Association for Computational Linguistics.