

YOUR FAIRNESS MAY VARY: PRETRAINED LANGUAGE MODEL FAIRNESS IN TOXIC TEXT CLASSIFICATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Warning: This paper contains samples of offensive text.

The popularity of pretrained language models in Natural Language Processing (NLP) systems calls for a careful evaluation of such models in down-stream tasks, which have a higher potential for societal impact. The evaluation of such systems usually focuses on *accuracy measures*. Our findings in this paper call for attention to be paid to *fairness measures* as well. Through the analysis of more than a dozen pretrained language models of varying sizes on two toxic text classification tasks, we demonstrate that focusing on accuracy measures alone can lead to models with wide variation in fairness characteristics. Specifically, we observe that fairness can vary even more than accuracy with increasing training data size and different random initializations. At the same time, we find that little of the fairness variation is explained by model size/compression, despite claims in the literature. To improve model fairness without retraining, we show that two post-processing methods developed for structured, tabular data can be successfully applied to a range of pretrained language models.

1 INTRODUCTION

Pre-trained, bidirectional language models (Devlin et al., 2019; Liu et al., 2019; Radford et al., 2019; Clark et al., 2020; He et al., 2021)¹ revolutionized natural language processing (NLP) research. LMs have provided a route to significant performance increases in several NLP tasks as demonstrated by the leaderboards (Rajpurkar et al., 2018; Wang et al., 2019a;b; AI2, 2021). More importantly, LMs have been applied to practical problems, leading to improved results for web search (Nayak, 2019) and have become an asset in fields such as medical evidence inference (Lehman et al., 2019; Subramanian et al., 2020) and even chemistry (Schwaller et al., 2021). These success stories have led to a sense in the community that application of pre-trained LMs to NLP tasks is beneficial.

While the progress in NLP tasks due to LMs is clear, the reasons behind this success are not as well understood (Rogers et al., 2021; McCoy et al., 2019), and there are also important downsides. In particular, several studies have documented the *bias* of LM-based models (Bolukbasi et al., 2016; Hutchinson et al., 2020; Webster et al., 2020; Borkan et al., 2019; de Vassimon Manela et al., 2021) and others discuss potential societal harms (Blodgett et al., 2020; Bender et al., 2021) for individuals or groups. We use the term *bias* to refer to systematic disparity in representation or outcomes for individuals based on their membership in certain protected groups such as gender, race, and ethnicity.

In this work, we focus on one important application of fine-tuned LMs, toxic text classification. Text toxicity predictors are already used in deployed systems (Perspective, 2021) and they are a crucial component for content moderation since online harassment is on the rise (Vogels, 2021). In down-stream applications such as toxic text classification, it is important to examine the behavior of LMs in terms of measures other than task-specific accuracy. This provides a more holistic understanding of model performance, leading to improved insights into appropriate uses of LMs for these tasks. As a first step toward this goal, we provide herein an empirical characterization of LMs for the task of toxic text classification using a combination of accuracy and bias measures, and study two post-processing methods for bias mitigation that have proved successful for structured, tabular data.

¹We use the acronym LM(s) to refer to language model(s) throughout the paper

One aspect of LMs that is hard to ignore is the increase in their size, as measured by the number of parameters in their architectures. In general, larger LMs seem to perform better on NLP tasks as they have the capacity to capture more complex correlations present in the training data. Bender et al. (2021) claim that this same property may also lead to more pronounced biases in their predictions. On the other hand, for *image* classification models that use large neural networks, Hooker et al. (2020) discuss how model pruning can lead to more biased predictions. In this work, we consider a wide variety of model architectures and sizes. We acknowledge that size is relative and what we consider large in this paper may not be considered as such in a different context.

We are interested in addressing the following questions regarding the effect of various factors on model performance:

1. *Model size*: Building on the work of Bender et al. (2021) and Hooker et al. (2020), how do the accuracy and group fairness of fine-tuned LM-based classifiers vary with their size?
2. *Random seeds*: LMs that start from different random initializations can behave differently in classification tasks. What is the effect of random seeds on the accuracy-fairness relationship?
3. *Data size*: The size of training/fine-tuning data is also an important dimension alongside model size. What happens to accuracy and fairness when more/less data is used for fine-tuning?
4. *Bias mitigation via post-processing*: Given the expense of training and fine-tuning large LMs, to what extent can we correct bias by only post-processing LM outputs?

We study the accuracy-fairness relationship in more than a dozen fine-tuned LMs for two different datasets that deal with prediction of text toxicity. The key contributions of our analysis are:

1. We empirically show that no blanket statement can be made regarding the fairness characteristics of fine-tuned LMs with respect to their size. It really depends on the combination of LM, task, and dataset.
2. We find that optimizing for accuracy measures alone can lead to models with wide variation in fairness characteristics. Specifically:
 - (a) While increasing data size for fine-tuning does not improve accuracy much beyond a point, the improvement in fairness is more significant and may continue after the improvement in accuracy has stopped for certain datasets and tasks. This suggests that choosing data size based on accuracy alone could lead to suboptimal performance with respect to fairness.
 - (b) While accuracy measures are known to vary (Dodge et al., 2020) with respect to different random initializations, the variation in fairness measures can be even greater.
3. We demonstrate that post-processing bias mitigation is an effective, computationally affordable solution to enhance fairness in fine-tuned LMs. In particular, one of the methods we experimented with allows for a large accuracy-fairness tradeoff space, leading to relative improvements of 50% for fairness, as measured by equalized odds, while reducing accuracy only by 2% (see Figure 5 religion group).

Our observations strengthen the chorus of recent work addressing bias mitigation in NLP in calling for a careful empirical analysis of fairness with fine-tuned LMs in the context of their application.

2 BACKGROUND, TERMINOLOGY AND RELATED WORK

Fairness in machine learning As machine learning models have become routinely deployed in practice, many studies noticed their tendency to perform unfairly in various contexts (Angwin et al., 2016; 2017; Buolamwini & Gebru, 2018; Bender et al., 2021; Park et al., 2021). To understand and measure model bias, researchers have proposed many definitions of algorithmic fairness. Broadly speaking, they fall into two categories: *group fairness* (Chouldechova & Roth, 2018) and *individual fairness* (Dwork et al., 2012). At a high level, group fairness requires similar average outcomes on different groups of individuals considered, for example comparable university acceptance rates

across ethnicities. Individual fairness requires similar outputs for similar individuals, e.g. two university applicants with similar credentials, but different ethnicity, gender, family background, etc., should either be both accepted or both rejected. In this paper we consider group fairness, noting that both have their pros and cons (Chouldechova & Roth, 2018; Dwork et al., 2012).

There are many definitions of group fairness and we refer to Verma & Rubin (2018) for a comprehensive overview. Statistical parity (SP) is one of the earlier definitions which requires the output of a model to be independent of the *protected attribute*, such as race or gender. In other words, the average outcome (e.g. prediction) across groups defined by the protected attribute needs to be similar. An alternative measure is equalized odds (EO) (Hardt et al., 2016), which requires the model output *conditioned* on the true label to be independent of the protected attribute. The violation of conditional independence for a given label (positive or negative) can be measured by the difference in accuracy across protected groups conditioned on that label. Taking the maximum or an average (average EO) of these label-specific differences quantifies the overall EO violation.

Many methods for achieving group fairness have been proposed. These methods are typically categorized as follows: (a) modifying the training data (pre-processing), (b) incorporating fairness constraints while training the model (in-processing), and (c) transforming the model output to enhance fairness (post-processing). A summary and implementation of group bias mitigation approaches are discussed in Bellamy et al. (2019). In this study, we investigate the use of post-processing methods to enhance fairness in classification tasks. We chose post-processing approaches since they do not require modification of training data or model training procedures, and hence can be efficiently applied to all LMs we consider. In addition, post-processing approaches could minimize the environmental impact of re-training/fine-tuning LMs (Bender et al., 2021; Patterson et al., 2021). We consider two post-processing approaches proposed by Wei et al. (2020) and Hardt et al. (2016), which have shown considerable success in mitigating bias for tabular data. Wei et al. (2020) optimize a score (predicted probability) transformation function to satisfy fairness constraints that are linear in conditional means of scores while minimizing a cross-entropy objective. Hardt et al. (2016) propose to solve a linear program to find probabilities with which to change the predicted output labels such that the equalized odds violation is minimized.

Fairness in Natural Language Processing In NLP systems, bias is broadly understood in two categories, intrinsic and extrinsic. Intrinsic bias refers to bias inherent in the representations, e.g. word embeddings used in NLP (Bolukbasi et al., 2016). Extrinsic bias refers to bias in downstream tasks, such as disparity in false positive rates across groups defined by protected attributes in a specified application. The concepts of intrinsic and extrinsic bias also correlate well with the notions of *representational* and *allocative* harms. While allocative harms correspond to disparities across the different groups in terms of decisions that lead to allocation of benefits/harms, representational harms are those perpetuated by representation of individuals in the feature space (Crawford, 2017). Abbasi et al. (2019) discuss how harms from stereotypical representations manifest as allocative harms later in the ML pipeline. However, probably because of their complexity, measuring intrinsic bias in the representations created by LMs may not necessarily reflect the behavior of models built by fine-tuning LMs. Goldfarb-Tarrant et al. (2021) discuss how intrinsic measures of bias do not correlate with extrinsic, application-specific, bias measures. Since we are concerned with the application of LMs to the specific task of toxic text classification, we restrict our focus to group fairness measures, which fall under the category of extrinsic bias. Previous work on bias mitigation in NLP has been focused on pre- and in-processing methods (Sun et al., 2019; Ball-Burack et al., 2021) and to the best of our knowledge, we are the first to use post-processing methods with NLP tasks.

3 METHODOLOGY

We are interested in studying how group fairness varies across different fine-tuned language models for binary classification. We choose to focus on text toxicity as the prediction task. Due to an increase in online harassment (Vogels, 2021) and the potential of both propagating harmful stereotypes of minority groups and/or inadvertently reducing their voices, the task of predicting toxicity in text has received increased attention in recent years (Kiritchenko et al., 2021). While we acknowledge that text toxicity presents different complex nuances (e.g., offensive text, harassment, hate speech), we focus on a binary task formulation. We adopt the definition of toxicity described in Borkan et al.

(2019) as *anything that is rude, disrespectful, or unreasonable that would make someone want to leave a conversation*.

Datasets We used two datasets that deal with toxic text classification: 1) Jigsaw, a large dataset released for a Kaggle competition (Jigsaw, 2019) that contains online comments on news articles, and 2) HateXplain, a dataset recently introduced with the intent of studying explanations for offensive and hate speech in Twitter and Twitter-like data (i.e., `gab.com`). Both datasets have fine-grained annotations for religion, race and gender. We used as protected groups the coarse-grained groups (e.g., religion) as opposed to the finer-grained annotations (e.g., Muslim, see below). Details about the sizes of the datasets, the splits we used and examples of text can be found in Appendix A.1.

Language models We consider more than a dozen LMs that cover a large spectrum of sizes. We selected the models to not only represent various sizes but also different styles of architecture and training. The models we included in our study are shown in Table 7 along with the number of parameters and the size of the *PyTorch* (Paszke et al., 2019) model on disk. If not specified, the version of the model used is *base*. For all our experiments, we used the Hugging Face implementation of Transformers (Wolf et al., 2020) and the corresponding implementations for all LMs in our study. In particular, we use the *text sequence classifier* without any modifications to increase reproducibility. For details on model sizes, fine-tuning and hyper-parameter tuning, as well as the computational infrastructure and costs, we refer the reader to Appendix A.2.

Fairness measures and implications In all our measurements, we considered two distinct groups: the protected group, consisting of text annotated as mentioning either religion, race or gender, and the rest of the samples, denoted as the unprotected group. Note we do not create fine-grained groups (e.g., Jewish), but consider larger groups (e.g., any mention of religion, such as Muslim, Jewish, atheist). We use equalized odds as the group fairness measure. Equalized odds is defined as the maximum of the true positive rate difference and false positive rate difference, where these differences are between the protected group and unprotected group. In toxic text classification, a true positive means that a toxic text is correctly identified as such, while a false positive means that a benign piece of text is marked as toxic. In terms of harms, a false negative (toxic text that is missed) may cause individuals to feel threatened or disrespected, while a false positive may be seen as censoring, which is particularly problematic if it reduces the voices of minority protected groups from online conversations. By using the protected groups of religion/race/gender mentioned above, we aim to reduce the effect of the presence or absence of religion/race/gender terms on the false negative and false positive rates. By taking the maximum, we are emphasizing the larger discrepancy as opposed to other studies that take the average of the two rate differences (average equalized odds). Note that unlike statistical parity, equalized odds does allow the protected (e.g., mention of religion) and unprotected (no religion) groups to have different toxicity (positive prediction) rates.

4 POST-PROCESSING METHODS FOR BIAS MITIGATION

We investigated the use of post-processing methods to mitigate violations of equalized odds. By post-processing, we mean methods that operate only on the outputs of the fine-tuned LMs and do not modify the models themselves.² The ability to avoid retraining models is a major advantage of post-processing due to the large computational cost of fine-tuning LMs. Post-processing also targets unfairness at a point closest to deployment and hence can have a direct impact on downstream operations that use the model predictions.

Hardt et al. (2016) The first post-processing method that we consider is by Hardt et al. (2016) (abbreviated HPS), who were the original proposers of the equalized odds criterion. We used the open-source implementation of their method from Bellamy et al. (2019), which post-processes binary predictions to satisfy EO while minimizing classification loss. While this method is effective in enforcing EO, one limitation is that it does not offer a trade-off between minimizing the deviation from EO and reducing the loss in accuracy.

Fair Score Transformer We also considered the Fair Score Transformer (FST) method of Wei et al. (2020), in part to provide the above-mentioned trade-off, and in part because it is a recent post-processing method shown to be competitive with several other methods (including in-processing).

²This is not to be confused with the post-processing of LM *embeddings*, before they are passed to classification layers. In this case, the classification layers must be retrained to account for the modified embeddings.

FST takes predicted probabilities (referred to as scores) as input and post-processes them to satisfy a fairness criterion. We choose *generalized equalized odds* (GEO), a score-based variant of EO, as the fairness criterion and then threshold the output score to produce a binary prediction. The application of FST required attention to three issues: 1) the provision of input scores that are indeed calibrated probabilities; 2) the choice of fairness parameter ϵ , which bounds the allowed GEO on the data used to fit FST; 3) the choice of binary classification threshold t . We consider a range of ϵ and t values to explore the trade-off between EO and accuracy. Due to numerical instability of the FST implementation in the original paper (occasional non-convergence in reasonable time for the Jigsaw dataset), we obtained a closed-form solution for one step in the optimization that leads to a more efficient implementation, running in minutes for all models and all datasets considered. More details on this implementation and the tuning of the parameters can be found in Appendix A.4.

Threshold post-processing We also tested the effect of thresholding alone, without fairness-enhancing transformation. We refer to this as *threshold post-processing* (TPP). This simple method corresponds to FST without calibrating the LM outputs, choosing ϵ large enough so that FST yields an identity transformation, and thresholding at level t .

5 THE ACCURACY-FAIRNESS RELATIONSHIP IN TOXIC TEXT CLASSIFICATION

We report on our study of the performance and fairness characteristics of several language models while varying parameters such as random seeds and the amount of training data. We also experiment with two post-processing methods for group bias mitigation and show that it is possible to reduce some of the bias presented by these models.

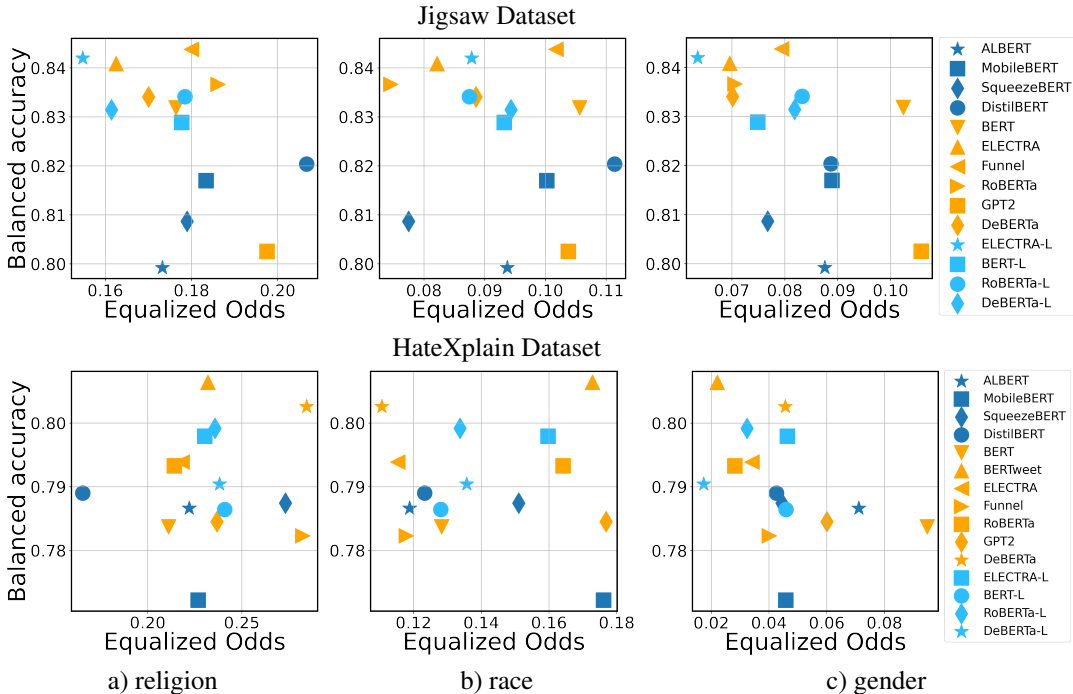


Figure 1: Balanced accuracy versus equalized odds for several fine-tuned LMs on the Jigsaw and HateXplain datasets.

5.1 CHARACTERIZATION OF LANGUAGE MODELS OF VARIED SIZES

The first set of experiments show how performance and fairness measures vary across models. In Figure 1 we show the performance as measured by balanced accuracy³ and the group fairness as

³We use balanced accuracy as a measure for performance as it is more honest, especially for the imbalanced Jigsaw dataset where a trivial predictor that always outputs the label “normal” would achieve $\sim 92\%$ accuracy.

measured by equalized odds on the x -axis (lower EO is better). The models are color-coded by their size - dark blue for small models, orange for regular size models and light blue for large models. From the figure, it can be noted that the variation in balanced accuracy is not as wide as the variation observed for equalized odds. For the HateXplain dataset, the gap between balanced accuracy and fairness variability is more prominent. In terms of accuracy (not balanced), the models perform even closer as shown in the plots in Appendix A.3. For EO, the spread is significant, with gaps of 0.10 between the largest and smallest values for Jigsaw, and 0.15 for HateXplain. Depending on the dataset and protected group, some larger models seem to lead to lower EO; for example, ELECTRA-large achieves best accuracy-EO results for religion as protected group (Jigsaw). For race, SqueezeBERT, which is one of the small models included in the study, achieves one of the best balanced accuracy-EO operating points (Jigsaw), hinting that size is not really correlated with the fairness of the model. Similarly, for HateXplain with religion as the protected group, DistilBERT, which is one of the small models, obtains the best balanced accuracy-EO operating point. In the next section, we will show models trained using various random seeds and show a low correlation between EO and model size.

These results strongly suggest that fairness measures should be included in the evaluation of language models. In the next sections, we show that, if fairness is not carefully considered, we can end up with models with widely varying fairness characteristics depending on the training conditions.

5.2 THE INFLUENCE OF RANDOM SEEDS ON ACCURACY AND FAIRNESS

Fine-tuning LMs depends on a random seed used for mini-batch sampling and for initializing the weights in the last layers of the network responsible for the binary classification. It is well documented in the literature that this random seed may influence the accuracy of the resulting model (Dodge et al., 2020). In Figure 2 we show that while balanced accuracy is somewhat stable, fairness can vary widely by only changing the random seed.

In fact, if we were to plot the accuracy instead of the balanced accuracy, all points would be virtually on a horizontal line for Jigsaw, as shown in Figure A.3. There are larger variations for EO. For Jigsaw, we observe a variation of up to 0.05 in equalized odds for some cases. In the case of HateXplain, the variation is considerably larger, with several models corresponding to a spread of 0.15 for the protected group of religion. The results in our experiments align with the ones shown in a recent study on underspecification in machine learning (D’Amour et al., 2020), where different random seeds lead to small variations in accuracy, but considerable variations in intrinsic bias as measured by gendered correlations.

To further probe whether there is a correlation between fairness and model size, we used the results for random seeds to compute the Pearson’s coefficient of correlation. These values are -0.357 for Jigsaw and -0.188 for HateXplain, with p-values of .000005 and 0.017, respectively. These results show a low correlation between fairness as measured by EO and model size.

5.3 LOW DATA REGIME

In general, it is known that more training data improves model accuracy. We experiment with training the models using a fraction of the training dataset, while keeping the test set the same. When the smaller datasets are subsampled from the original dataset, we ensure that the larger datasets include the smaller ones to simulate situations when more data is collected and used for training. The results are shown for one small/regular/large model in Figure 3. Each data point in the graph represents the average of eleven runs performed with different random seeds, one for each run. In very few cases, the random seed led to a degenerate model and we did not include these runs in the averaged results. Overall, there were about five degenerate runs for each dataset.

We observe that in the case of Jigsaw, equalized odds generally keeps improving even when the accuracy plateaus, suggesting that, from a fairness point of view, it may be beneficial to collect more data for training. This does not seem to be the case for the HateXplain dataset, where the accuracy does not plateau and the fairness measure oscillates. One of the reasons could be that HateXplain is much smaller in size than Jigsaw and hence Jigsaw’s training is much more stable. Similar trends are observed for the rest of the models in our study.

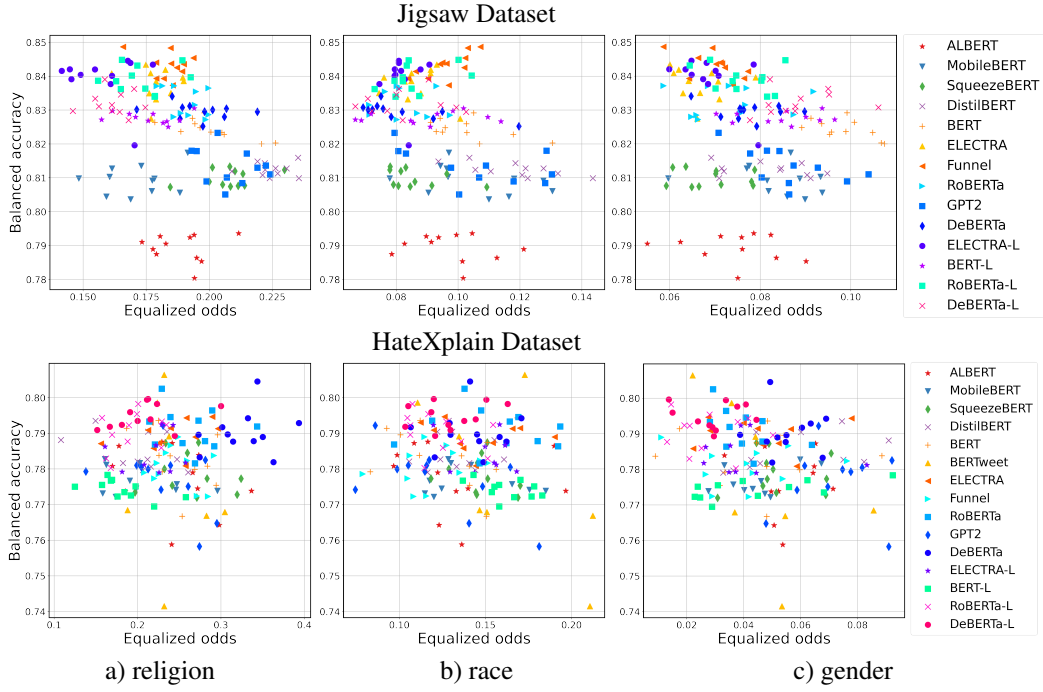


Figure 2: Balanced accuracy versus equalized odds for several fine-tuned LMs when varying only the random seed used in fine-tuning.

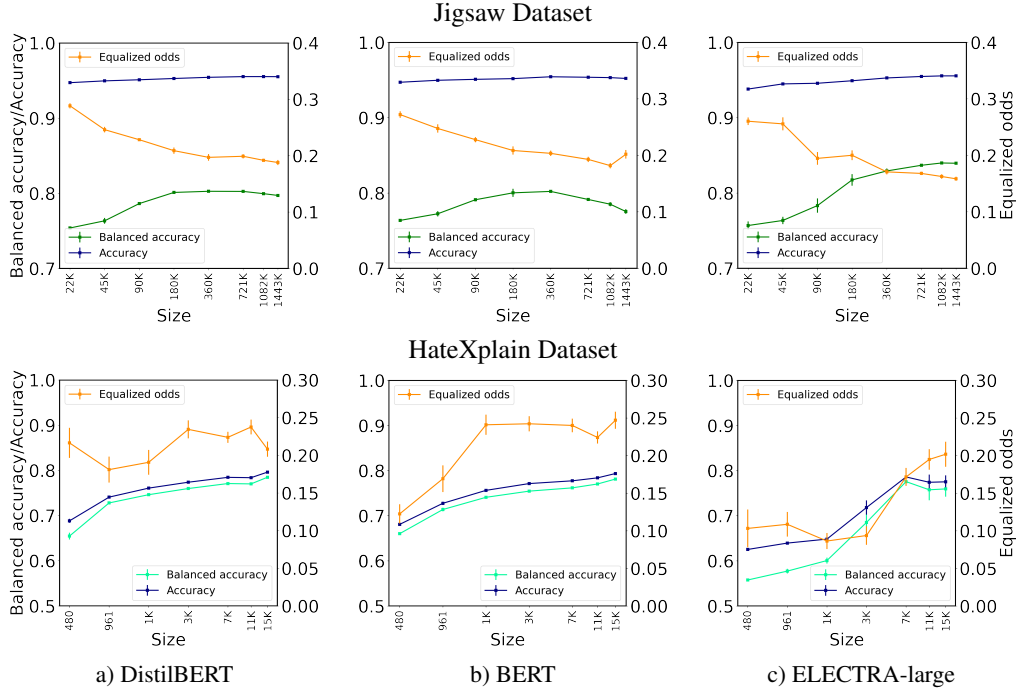


Figure 3: Accuracy, balanced accuracy and equalized odds (religion) for fine-tuned LMs when varying the amount of data used in training and the random seeds. Error bars denote the ± 1 SE (standard error) of the mean.

5.4 BIAS MITIGATION THROUGH POST-PROCESSING METHODS

In this section we experiment with applying post-processing methods for group bias mitigation. We first discuss the results of parameter tuning for Fair Score Transformer (FST) (Wei et al., 2020). More details can be found in the appendix A.4. The FST method has one parameter, ϵ , that can be fine-tuned. Using the transformed scores from the FST, we also investigate tuning the threshold used in the binary classifier, instead of using the default value of 0.5, as explained in Section 4. Figure 4 depicts the data points obtained by varying epsilon and for each epsilon value, varying the classification threshold. ⁴ When choosing an operating point, the points on the black Pareto frontier are the most interesting points: highest balanced accuracy and lowest equalized odds. For reference, we also show the baseline points without bias mitigation for the dev and test sets. All data points are plotted for fine-tuned BERT. Similar trends are observed for the rest of the models considered in this study and for the HateXplain dataset.




Figure 4: BERT: Balanced accuracy versus equalized odds for several fine-tuned LMs on the Jigsaw dataset when varying epsilon and the threshold for binary classification after applying the FST method for group bias mitigation.

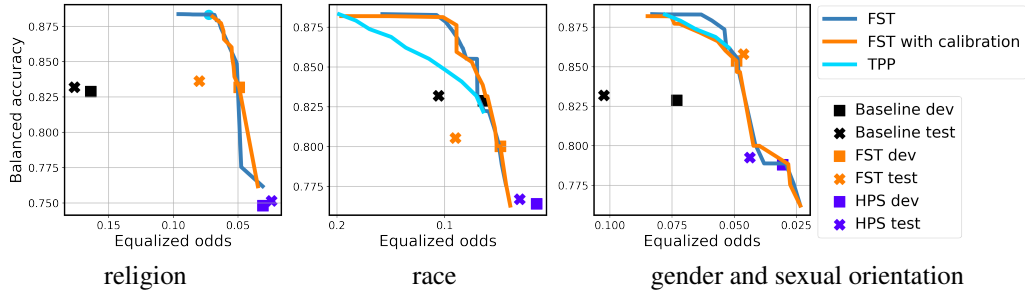


Figure 5: BERT: Balanced accuracy versus equalized odds on the Jigsaw dataset when applying the FST and HPS methods for group bias mitigation and threshold post-processing (TPP) alone.

We also experimented with calibrating the scores using logistic regression before post-processing. In Figure 5, we plot the Pareto frontiers of bias mitigation when applying FST, with and without calibration, along with the threshold post-processing (TPP) method. We also show the result of HPS, which yields a single operating point, as well as the baselines without bias mitigation. In general on the Jigsaw dataset, FST is successful in reducing EO with different degrees of success depending on the model/group. It thus offers an interesting set of points with different accuracy-EO trade-offs. For reference, we show the equivalent point for the test set (orange x) for the operating point in dev that achieves an equalized odds of at most 0.05 (orange square). In certain cases, FST manages to lower the equalized odds with minimal or no decrease in accuracy, as seen in the religion and gender columns in Figure 5. Note that all points in the plots except for the x points are plotted using the dev dataset split, the x points are test points corresponding to dev points that obtain an EO of at most 0.05.

⁴ All points are shown for the dev set as this plot corresponds to hyper-tuning FST parameters.




Figure 6: Balanced accuracy versus equalized odds for xseveral fine-tuned LMs (religion group) on the HateXplain dataset when applying the FST and HPS methods for group bias mitigation and threshold post-processing (TPP) alone.

In comparison, HPS seems particularly effective in lowering the equalized odds and thus improving the fairness of the model, with some penalty on the accuracy. For Jigsaw, applying only TPP (i.e., tuning the threshold used in the binary classification) also offers some interesting operating points. TPP has a small search space compared to FST and sometimes the Pareto frontier is reduced to one point, as is the case for the religion group. In general, FST has superior Pareto frontiers compared to TPP alone. In addition, as we will discuss shortly, TPP proved inefficient for the HateXplain dataset. Last, using score calibration before feeding the scores to FST does not seem to offer significant improvements. Similar trends can be observed for the rest of the models.

In Figure 6, we show the results of applying bias mitigation techniques for a few LMs, one for each size category, on the HateXplain dataset with religion as the protected group. Unlike Jigsaw, the results of the bias mitigation techniques follow different trends. HPS still manages to substantially reduce the EO for all models, but with a considerable decrease in balanced accuracy (in some cases, more than six percentage points). For FST, the fine-tuning for epsilon and classification threshold does not lead to a large search space as observed in the Jigsaw case. Moreover, the reduction in EO is more limited and sometimes the improvement observed for the dev set disappears in test. There are cases, though, such as BERT, where FST successfully reduces EO and the reduction is maintained or even improved in test. Across the board, tuning only the threshold used in classification (TPP) did not lead to improved results and we omit showing them in the plots.

Overall, we find the post-processing methods for bias mitigation worth considering. They are straightforward to apply, run in the order of seconds or minutes on the CPU of a laptop and they offer interesting operating points when other methods for bias elimination would incur a significant computational cost, such as pre-processing or in-processing techniques. Obtaining the Pareto frontiers is instantaneous as the search space for FST is not that large.

6 LIMITATIONS AND CONCLUSIONS

We presented a comprehensive study of language models and their performance/fairness relationship. We chose several models to cover different sizes and different architectures. While we did not consider some of the largest recent models available, we believe we have experimented with a wide variety of models that have been discussed well in the literature. Using A100 GPUs, we were able to finish fine-tuning for our largest models in at most 24 hours. One important aspect we would like to emphasize is that identifying toxic text is not an easy task, not even for humans. As such, we expect the datasets to be noisy and contain samples that are not annotated correctly. Upon manual inspection, we could identify some samples for which we did not agree with their labels. As a consequence, while we expect the trends shown in this paper to hold, the actual absolute numbers may vary with datasets. We hope that this study can drive the following point across: we cannot make a blanket statement on the fairness of language models with respect to their size or architecture, while training factors such as data size and random seeds can make a large difference. This makes it all the more important for researchers/practitioners to make fairness an integral part of the performance evaluation of language models before deployment.

7 REPRODUCIBILITY STATEMENT

The data processing we performed for the datasets we used is briefly explained in Appendix A.1. In all our experiments we used unmodified versions of the model implementations from the Hugging Face transformers library and the main scripts to tune the models are modified versions of the sequence text classification examples accompanying the library. The hyper-parameter tuning we performed was minimal (varying the epochs size 1-3, two values for learning rates, 11 values for random seeds). More details on the experimental infrastructure can be found in Appendix A.2. The limiting factor in reproducing these results is having access to GPUs such as the V100 and A100. We could not find a public implementation for FST. We provide details on our implementation in Appendix A.4. HPS is open source and can be found at https://github.com/Trusted-AI/AIF360/blob/master/aif360/algorithms/postprocessing/eq_odds_postprocessing.py.


8 ETHICS STATEMENT

In this work, we attempted to address the following research questions for language models: how do model size, training size, random seeds affect the relationship between performance and fairness (as measured by equalized odds)? Can post-processing methods for bias mitigation lead to better operating points for both accuracy and fairness? We find these questions important to ask in the context of the ethics of using language models in text toxicity prediction, in particular, and in NLP research, in general. This study used a considerable amount of computational resources and this is our main ethics concern for conducting this study. We did try to keep the number and the size of models we experimented with limited to reduce the carbon footprint of the study. We hope the results we show in this paper are worth the computational resources used to perform the experiments.

Bias mitigation can lead to undesirable outcomes. For example, one aspect we did not look into is what happens with other groups when the mitigation is applied only for one of the groups. Similarly, we did not consider intersectionality. We focused only on group fairness and do not provide any insights into individual fairness. Last, but not least, abstract metrics have limitations and the societal impacts resulting from bias mitigation are not well understood (Olteanu et al., 2017). These issues are universal to bias mitigation techniques and not particular to our use case.

REFERENCES

- Mohsen Abbasi, Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. Fairness in representation: quantifying stereotyping as a representational harm. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pp. 801–809. SIAM, 2019.
- Allen Institute for AI AI2. Leaderboards, 2021. URL <https://leaderboard.allenai.org/>.
- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing, May 2016.
- Julia Angwin, Jeff Larson, Lauren Kirchner, and Surya Mattu. Minority Neighborhoods Pay Higher Car Insurance Premiums Than White Areas With the Same Risk. <https://www.propublica.org/article/minority-neighborhoods-higher-car-insurance-premiums-white-areas-same-risk>, April 2017.
- Ari Ball-Burack, Michelle Seng Ah Lee, Jennifer Cobbe, and Jatinder Singh. Differential tweetment: Mitigating racial dialect bias in harmful tweet detection. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 116–128, 2021.
- Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, et al. Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5):4–1, 2019.

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? . In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, pp. 610–623, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383097. doi: 10.1145/3442188.3445922. URL <https://doi.org/10.1145/3442188.3445922>.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. Language (technology) is power: A critical survey of “bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5454–5476, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.485. URL <https://www.aclweb.org/anthology/2020.acl-main.485>.
- Aja Bogdanoff. Saying goodbye to Civil Comments. https://medium.com/@aja_15265/saying-goodbye-to-civil-comments-41859d3a2b1d, 2017. [Online; accessed 21-July-2021].
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 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*, pp. 4356–4364, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.
- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion of The 2019 World Wide Web Conference, WWW, 2019*. doi: 10.1145/3308560.3317593. URL <https://doi.org/10.1145/3308560.3317593>.
- Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pp. 77–91. PMLR, 2018.
- Alexandra Chouldechova and Aaron Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.
- Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. Leveraging labeled and unlabeled data for consistent fair binary classification. *Advances in Neural Information Processing Systems*, 32:12760–12770, 2019.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. ELECTRA: pre-training text encoders as discriminators rather than generators. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=r1xMH1BtvB>.
- Kate Crawford. The trouble with bias. https://www.youtube.com/watch?v=fMym_BKWQzk, December 2017.
- Zihang Dai, Guokun Lai, Yiming Yang, and Quoc Le. Funnel-transformer: Filtering out sequential redundancy for efficient language processing. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/2cd2915e69546904e4e5d4a2ac9e1652-Abstract.html>.
- Alexander D’Amour, Katherine A. Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yi-An Ma, Cory Y. McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. Underspecification presents challenges for credibility in modern machine learning. *CoRR*, abs/2011.03395, 2020. URL <https://arxiv.org/abs/2011.03395>.

- Daniel de Vassimon Manela, David Errington, Thomas Fisher, Boris van Breugel, and Pasquale Minervini. Stereotype and skew: Quantifying gender bias in pre-trained and fine-tuned language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2232–2242, Online, April 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.eacl-main.190>.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah A. Smith. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *CoRR*, abs/2002.06305, 2020. URL <https://arxiv.org/abs/2002.06305>.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226, 2012.
- Benjamin Fish, Jeremy Kun, and Ádám D Lelkes. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pp. 144–152. SIAM, 2016.
- Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sanchez, Mugdha Pandya, and Adam Lopez. Intrinsic bias metrics do not correlate with application bias. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, Online, 2021. Association for Computational Linguistics. URL <https://arxiv.org/abs/2012.15859>.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 2017. URL <http://proceedings.mlr.press/v70/guo17a.html>.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29:3315–3323, 2016.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: decoding-enhanced BERT with disentangled attention. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=XPZJaotutsD>.
- Sara Hooker, Nyalleng Moorosi, Gregory Clark, S. Bengio, and Emily L. Denton. Characterising bias in compressed models. *ArXiv*, abs/2010.03058, 2020. URL <https://arxiv.org/abs/2010.03058>.
- Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. Social biases in NLP models as barriers for persons with disabilities. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5491–5501, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.487. URL <https://aclanthology.org/2020.acl-main.487>.
- Forrest Iandola, Albert Shaw, Ravi Krishna, and Kurt Keutzer. SqueezeBERT: What can computer vision teach NLP about efficient neural networks? In *Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing*, pp. 124–135, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.sustainlp-1.17. URL <https://aclanthology.org/2020.sustainlp-1.17>.
- Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. Wasserstein fair classification. In *Uncertainty in Artificial Intelligence*, pp. 862–872. PMLR, 2020.
- Kaggle Jigsaw. Jigsaw Unintended Bias in Toxicity Classification. <https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>, 2019. [Online; accessed 21-July-2021].

- Faisal Kamiran, Asim Karim, and Xiangliang Zhang. Decision theory for discrimination-aware classification. In *2012 IEEE 12th International Conference on Data Mining*, pp. 924–929. IEEE, 2012.
- Michael P Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 247–254, 2019.
- Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C. Fraser. Confronting abusive language online: A survey from the ethical and human rights perspective. *Journal of Artificial Intelligence Research*, 2021.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=H1eA7AEtvS>.
- Eric Lehman, Jay DeYoung, Regina Barzilay, and Byron C Wallace. Inferring which medical treatments work from reports of clinical trials. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pp. 3705–3717, 2019.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. Hateexplain: A benchmark dataset for explainable hate speech detection. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 14867–14875. AAAI Press, 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17745>.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3428–3448, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1334. URL <https://aclanthology.org/P19-1334>.
- Pandu Nayak. Understanding searches better than ever before, Oct 2019. URL <https://blog.google/products/search/search-language-understanding-bert/>.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. Bertweet: A pre-trained language model for english tweets. In Qun Liu and David Schlangen (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pp. 9–14. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.emnlp-demos.2. URL <https://doi.org/10.18653/v1/2020.emnlp-demos.2>.
- Alexandra Olteanu, Kartik Talamadupula, and Kush R. Varshney. The limits of abstract evaluation metrics: The case of hate speech detection. In Peter Fox, Deborah L. McGuinness, Lindsay Poirier, Paolo Boldi, and Katharina Kinder-Kurlanda (eds.), *Proceedings of the 2017 ACM on Web Science Conference, WebSci 2017, Troy, NY, USA, June 25 - 28, 2017*, pp. 405–406. ACM, 2017. doi: 10.1145/3091478.3098871. URL <https://doi.org/10.1145/3091478.3098871>.
- Yoonyoung Park, Jianying Hu, Moninder Singh, Issa Sylla, Irene Dankwa-Mullan, Eileen Koski, and Amar K. Das. Comparison of Methods to Reduce Bias From Clinical Prediction Models of Postpartum Depression. *JAMA Network Open*, 4(4):e213909–e213909, 04 2021. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2021.3909. URL <https://doi.org/10.1001/jamanetworkopen.2021.3909>.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021.
- Perspective. Using Machine Learning to Reduce Toxicity Online. <https://perspectiveapi.com/how-it-works/>, 2021. [Online; accessed 21-July-2021].
- Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. *arXiv preprint arXiv:1709.02012*, 2017.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are Unsupervised Multitask Learners. 2019. URL <https://openai.com/blog/better-language-models/>.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don’t know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 784–789, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2124. URL <https://aclanthology.org/P18-2124>.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2021.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2020.
- Philippe Schwaller, Daniel Probst, Alain C. Vaucher, Vishnu H. Nair, David Kreutter, Teodoro Laino, and Jean-Louis Reymond. Mapping the space of chemical reactions using attention-based neural networks. *Nature Machine Intelligence*, 3(2):144–152, Feb 2021. ISSN 2522-5839. doi: 10.1038/s42256-020-00284-w. URL <https://doi.org/10.1038/s42256-020-00284-w>.
- Shivashankar Subramanian, Ioana Baldini, Sushma Ravichandran, Dmitriy Katz-Rogozhnikov, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Varshney Kush R, Annmarie Wang, Pradeep Mangalath, and Laura Kleiman. A natural language processing system for extracting evidence of drug repurposing from scientific publications. *Proceedings of the AAAI Conference on Innovative Applications of Artificial Intelligence*, 2020.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. Mitigating gender bias in natural language processing: Literature review. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1630–1640, 2019.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. Mobilebert: a compact task-agnostic BERT for resource-limited devices. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 2158–2170. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.195. URL <https://doi.org/10.18653/v1/2020.acl-main.195>.
- Sahil Verma and Julia Rubin. Fairness definitions explained. In *Proceedings of the International Workshop on Software Fairness, FairWare ’18*, pp. 1–7, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450357463. doi: 10.1145/3194770.3194776. URL <https://doi.org/10.1145/3194770.3194776>.

- Emily A. Vogels. The State of Online Harassment. <https://www.pewresearch.org/internet/2021/01/13/the-state-of-online-harassment/>, 2021. [Online; accessed 21-July-2021].
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. *arXiv preprint 1905.00537*, 2019a.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. 2019b. In the Proceedings of ICLR.
- Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, and Slav Petrov. Measuring and reducing gendered correlations in pre-trained models. *CoRR*, abs/2010.06032, 2020. URL <https://arxiv.org/abs/2010.06032>.
- Dennis Wei, Karthikeyan Natesan Ramamurthy, and Flávio du Pin Calmon. Optimized score transformation for fair classification. In Silvia Chiappa and Roberto Calandra (eds.), *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1673–1683. PMLR, 2020. URL <http://proceedings.mlr.press/v108/wei20a.html>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- Blake Woodworth, Suriya Gunasekar, Mesrob I Ohannessian, and Nathan Srebro. Learning non-discriminatory predictors. In *Conference on Learning Theory*, pp. 1920–1953. PMLR, 2017.
- Forest Yang, Mouhamadou Cisse, and Oluwasanmi O Koyejo. Fairness with overlapping groups; a probabilistic perspective. *Advances in Neural Information Processing Systems*, 33, 2020.

A APPENDIX

In this section, we discuss all aspects related to the methodology of the experiments we performed, include additional experimental results and provide more details on post-processing methods for bias mitigation.

A.1 DATASETS

A.1.1 JIGSAW UNINTENDED BIAS IN TEXT TOXICITY CLASSIFICATION

In 2019, Jigsaw released a large dataset as part of a public Kaggle competition (Jigsaw, 2019). The dataset is a collection of roughly two million samples of text from online discussions (Bogdanoff, 2017). The samples are rated for toxicity and annotated with attributes for protected groups. Table 1 shows the groups we considered in our analysis and the available fine-grained group annotations. Note that we considered the coarser groups; a sample text belongs to a protected (coarse) group if any (fine-grained) annotation for the sample text exists. We used the original training dataset split in a 80/20 ratio for training and development (dev) tuning, respectively. For reporting test results, we used the private test split released on Kaggle. Statistics for the dataset splits are shown in Table 3. Each sample in the dataset (see Table 2 for a few samples from the dataset) has a toxicity score and we consider anything higher than 0.5 to be toxic.

Table 1: The identity groups for Jigsaw dataset with their corresponding fine-grained annotations.

Group	Fine-grained annotation
religion	atheist, buddhist, christian, hindu, jewish, other religion
race	white, asian, black, latino, other race or ethnicity
gender and sexual orientation	bisexual, female, male, heterosexual, homosexual gay or lesbian, transgender, other gender, other sexual orientation

Table 2: Jigsaw dataset samples

Comment text	Toxicity	Group
The Atwood fable is Donald, is it? My impression of this noise (over Atwood) is that it's a gimmick by Atwood and her publisher to cash in on the Donald effect. As if we needed slaves in bonnets to remind us that Donald is a jerk (and where was Atwood's novel when Monica was being pawed over?). A word to defenders of women: don't spend your political capital on stupid analogies.	Toxic	Gender
I got a question for you, dear, and it is a fair question: We all know what is happening in Syria; where are all the women's marches over the slaughter in that country?. And, why has Trudeau been silent, like his pal Barry Obama, on taking effective military action against Syria? All you lefties are the same: you have no side vision.	Normal	Gender

A.1.2 HATEXPLAIN: TOXIC TEXT IN TWITTER AND TWITTER-LIKE TEXT

HateXplain (Mathew et al., 2021) was recently introduced with the intent of studying explanations in offensive and hate speech in Twitter and Twitter like data (i.e., `gab.com`). For the purposes of our study, we collapse the annotations for offensive and hate speech into one class of toxic text. Similar to the Jigsaw dataset, HateXplain samples have fine-grained annotations for protected groups. We use as groups the coarse-level annotations, as we did for the Jigsaw dataset. The groups that we consider are presented in Table 4 and a few examples from the dataset are shown in Table 5. Note the text in each sample is represented in the dataset as a list of tokens; in the table, we concatenated them with spaces and this is the way we use them as inputs for the classifiers as well. We used the splits as provided in the dataset; dataset statistics are shown in Table 6.

A.2 LANGUAGE MODELS AND COMPUTATION INFRASTRUCTURE USED IN TUNING

We consider more than a dozen language models that cover a large spectrum of sizes. We selected the models covering not only various sizes, but also different styles of architecture and training. The models we included in our study are shown in Table 7 along with the number of parameters and the size of the *PyTorch* Paszke et al. (2019) model on disk. If not specified, the version of the model used is *base*.

For all our experiments, we used the Hugging Face implementation of Transformers Wolf et al. (2020) and the corresponding implementations for all the language models used in our study. In particular, we use the *text sequence classifier* without any modifications to increase the reproducibility of this study. For most experiments, we run model fine-tuning for 1-3 epochs and choose the best model based on the highest accuracy obtained on the dev split. When presenting experimental re-

Table 3: Jigsaw dataset statistics: sample counts per dataset split and protected group

Dataset split	Total	Religion	Race	Gender
Train	1443899	50748	31022	70703
Dev	360975	12769	7999	17869
Test	97320	3316	1911	4367

Table 4: The identity groups for HateXplain dataset with their corresponding fine-grained annotations.

Group	Fine-grained annotation
race	African, Arab, Asian, Caucasian, Hispanic
religion	Islam, Buddhism, Jewish, Hindu, Christian
gender	Men, Women

Table 5: HateXplain dataset samples

Text	Toxicity	Group
why is anna acting like she black	Normal	Race
like brazil many are descendants of west africa social constructs labels classifications have ruined a many people who falsely self identify as white or just do not even know abt their native land all b c of them bum ass cauczoids aka caucasians	Toxic	Race

sults, we focus on balanced accuracy as the Jigsaw dataset is imbalanced and looking only at accuracy results may be misleading. In general, higher accuracy leads to higher balanced accuracy, with the exception of two models – GPT2 and SqueezeBERT. For these two, the best balanced accuracy is less than 2 percentage points higher than the balanced accuracy corresponding to the highest overall accuracy across the various hyper-parameter runs. We also experiment with two learning rates and observe that the large models tend to prefer smaller learning rates, degenerating for lower learning rates. For large models with Jigsaw we fine-tune only for one epoch to keep the compute time under 24 hours. The model accuracy we obtained are in line with state-of-the-art results for these types of tasks.

The large models are fine-tuned using the A100 Nvidia GPUs, while the rest of the models are fine-tuned on V100 Nvidia GPUs. The experiments for HateXplain dataset run from 10 minutes to under an hour, while the experiments for the large models with the Jigsaw dataset can take up to 24 hours.

A.3 THE INFLUENCE OF RANDOM SEEDS ON ACCURACY AND EQUALIZED ODDS

In this section we present graphs similar to the ones in Section 5.2 using accuracy as a measure of performance instead of balanced accuracy. These plots makes it obvious how close in performance all models are and emphasize the gap in fairness measure observed across different random seeds for each fine-tuned model. The results are shown in Figure 7. Note that all Jigsaw models get an accuracy in performance of approximately 95% with a gap of approximately .05 for equalized odds. HateXplain models exhibit a higher variance in accuracy (4-5%) with an even larger gap of .15 for equalized odds for most models.

A.4 FAIR SCORE TRANSFORMER (FST)

FST takes predicted probabilities (referred to as scores) as input and outputs scores satisfying a fairness criterion. We choose *generalized equalized odds* (GEO), a score-based variant of EO, as the fairness criterion and then threshold the output score to produce a binary prediction. Instead of races, GEO is computed as the maximum between the group-wise difference in the scores for the true

Table 6: HateXplain dataset statistics: sample counts per dataset split and protected group

Dataset split	Total	Religion	Race	Gender
Train	15383	3924	5418	3102
Dev	1922	481	672	396
Test	1924	468	685	375

Table 7: The size (number of parameters, size on disk) for the language models considered in this study.

Group Size	Language Model	# of parameters	Size on disk
Small	ALBERT (Lan et al., 2020)	12M	45MB
	MobileBERT (Sun et al., 2020)	25.3M	95MB
	SqueezeBERT (Iandola et al., 2020)	51M*	196MB
	DistilBERT (Sanh et al., 2020)	66M	256MB
Regular	BERT (Devlin et al., 2019)	110M	418MB
	BERTweet (Nguyen et al., 2020)	110M	418MB
	ELECTRA (Clark et al., 2020)	110M	418MB
	Funnel (small) (Dai et al., 2020)	117M*	444MB
	RoBERTa (Liu et al., 2019)	125M	476MB
	GPT2 (Radford et al., 2019)	117M	487MB
	DeBERTa (He et al., 2021)	140M	532MB
Large	ELECTRA-large	335M	1.3GB
	BERT-large	340M	1.3GB
	RoBERTa-large	355M	1.4GB
	DeBERTa-large	400M	1.6GB

*Approximate number of parameters, as exact parameter size could not be found.

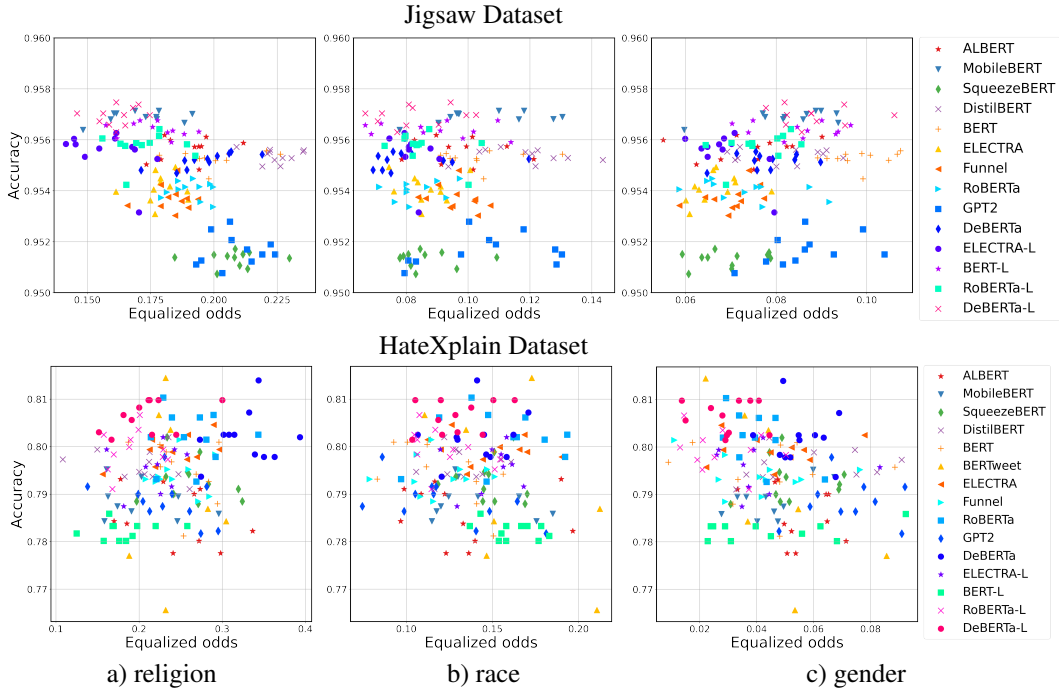


Figure 7: Accuracy versus equalized odds for several fine-tuned LMs when varying only the random seed used in fine-tuning.

positive and the false positive, respectively, where instead of the predicted label, the corresponding probability for the label is used instead.

The application of FST requires attention to three issues: 1) the provision of input scores that are indeed calibrated probabilities; 2) the choice of fairness parameter ϵ , which bounds the allowed GEO on the data used to fit FST; 3) the choice of binary classification threshold t .

Regarding issue 1), we found that the distributions of softmax outputs of the tested LMs are bi-modal and highly concentrated near values of 0 and 1 (as commonly observed with deep neural

networks). Such skewed distributions appear to violate FST’s assumption of probabilities as input and are typically not encountered on tabular datasets on which FST was previously tested. Thus we experimented with calibrating the LM outputs. We considered both logistic regression of the class label on the logit outputs of the LMs (a generalization of temperature scaling (Guo et al., 2017)), as well as linear regression on the logit outputs followed by clipping of the resulting values to the interval $[0, 1]$. In general, logistic regression proved somewhat beneficial for the Jigsaw dataset and we included it in our results.

Regarding issue 2), we found, as noted by Wei et al. (2020), that while the parameter ϵ controls the deviation from GEO (i.e. the “GEO difference”), this is not always correlated with the EO difference, which is a function of the output after thresholding. Regarding 3), we found that varying the threshold t can significantly affect equalized odds as well as accuracy and balanced accuracy, and can sometimes even produce a reasonable trade-off between them. For this reason, we included a version of post-processing (see “Threshold post-processing” in Section 4. This effect of the prediction threshold on fairness has not been explored in previous work to our knowledge.

As a result of our observations regarding 2) and 3), we used the following procedure to select a *set* of (ϵ, t) pairs to map out a trade-off between fairness and performance. The training set used to fine-tune the LMs is never seen by FST. The development dataset (“dev”) is used to both tune the FST parameters and evaluate the resulting transformation. As such, the dev dataset was further split into a dev-train set and a dev-eval set. Given an ϵ value, FST was fit on the dev-train set to ensure a GEO difference of at most ϵ . Then on the dev-eval set, given ϵ and t , scores were transformed by FST with parameter ϵ , thresholded at level t to produce a binary label, and finally evaluated for both fairness and performance. Each (ϵ, t) pair thus yields one point in the equalized odds-performance plane, as seen in Figure 4. We selected (ϵ, t) pairs that are Pareto-efficient on the dev-eval set, to ensure the best fairness-accuracy trade-off.

This is the first time FST is used with unstructured, text data and with large datasets in the order of millions of samples. First, we implemented FST following the proposed implementation in Wei et al. (2020). This first implementation ended up with numerical instabilities that lead to either slow running times (in the order of hours) or even situation when the method did not converge. We managed to improve upon the computational cost of FST, which was instrumental in scaling to the large Jigsaw dataset and allowing rapid experimentation. Specifically, in the dual ADMM algorithm of Wei et al. (2020), the first step (eq. (14) therein) consists of n parallel optimizations, each involving a single variable. We observed that these optimizations can be done in closed form by solving a cubic equation. The replacement of an iterative optimization with a closed-form solution greatly reduces the computational cost of FST. FST runs in the order of 1-2 minutes for the Jigsaw dataset and in seconds for HateXplain. Equally important, it also eliminates instances of the iterative optimization failing to converge.

A.5 POST-PROCESSING METHODS FOR BIAS MITIGATION

In addition to the two post-processing methods that we considered in our study, other post-processing methods for bias mitigation include assigning favorable labels to unprivileged groups in regions of high classifier uncertainty (Kamiran et al., 2012), minimizing error disparity while maintaining classifier calibration (Pleiss et al., 2017), a relaxed nearly-optimal procedure for optimizing equalized odds (Woodworth et al., 2017), shifting the decision boundary for the protected group (Fish et al., 2016), iterative post-processing to achieve unbiased predictions on every identifiable subpopulation (Kim et al., 2019), recalibrating a classifier using a group-dependent threshold to optimize equality of opportunity (defined as the difference between the group-wise true positive rates) (Chzhen et al., 2019), using optimal transport to ensure similarity in group-wise predicted score distributions (Jiang et al., 2020), and a plug-in approach for transforming the predicted probabilities to satisfy fairness constraints (Yang et al., 2020).

A.6 MATHEMATICAL DEFINITIONS FOR THE MEASURES USED

Accuracy

$$accuracy = \frac{\text{number of samples predicted correctly}}{\text{total number of samples}}$$

Balanced accuracy

$$\text{balanced accuracy} = \frac{\frac{\text{true negatives}}{\text{total number of negative samples}} + \frac{\text{true positives}}{\text{total number of positive samples}}}{2}$$

Equalized odds

$$TPR = \frac{\text{number of samples predicted correctly as positive}}{\text{total number of positive samples}}$$

$$FPR = \frac{\text{number of samples predicted incorrectly as positive}}{\text{total number of negative samples}}$$

$$EO = \max(\text{abs}(TPR_{\text{protected group}} - TPR_{\text{unprotected group}}), \text{abs}(FPR_{\text{protected group}} - FPR_{\text{unprotected group}}))$$

Generalized equalized odds

$$\text{generalized TPR} = \text{genTPR} = \frac{\text{sum of predicted probabilities for positive samples}}{\text{total number of positive samples}}$$

$$\text{generalized TNR} = \text{genTNR} = \frac{\text{sum of predicted probabilities for negative samples}}{\text{total number of negative samples}}$$

$$\text{generalized EO} = \max(\text{abs}(\text{genTPR}_{\text{protected}} - \text{genTPR}_{\text{unprotected}}), \text{abs}(\text{genTNR}_{\text{protected}} - \text{genTNR}_{\text{unprotected}}))$$