

Demographics Should Not Be the Reason of Toxicity: Mitigating Discrimination in Text Classifications with Instance Weighting

Guanhua Zhang^{†§*}, Bing Bai^{§*}, Junqi Zhang[§], Kun Bai[§], Conghui Zhu[‡], Tiejun Zhao[‡]

[‡]Harbin Institute of Technology, [§]Tencent CSIG

{guanhzhang, icebai, benjqzhang, kunbai}@tencent.com,
{chzhu, tjzhao}@hit-mlab.net

Abstract

With the recent proliferation of the use of text classifications, researchers have found that there are certain unintended biases in text classification datasets. For example, texts containing some demographic identity-terms (*e.g.*, “gay”, “black”) are more likely to be abusive in existing abusive language detection datasets. As a result, models trained with these datasets may consider sentences like “She makes me happy to be gay” as abusive simply because of the word “gay.” In this paper, we formalize the unintended biases in text classification datasets as a kind of *selection bias* from the *non-discrimination distribution* to the *discrimination distribution*. Based on this formalization, we further propose a model-agnostic debiasing training framework by recovering the non-discrimination distribution using instance weighting, which does not require any extra resources or annotations apart from a pre-defined set of demographic identity-terms. Experiments demonstrate that our method can effectively alleviate the impacts of the unintended biases without significantly hurting models’ generalization ability.

1 Introduction

With the development of Natural Language Processing (NLP) techniques, Machine Learning (ML) models are being applied in continuously expanding areas (*e.g.*, to detect spam emails, to filter resumes, to detect abusive comments), and they are affecting everybody’s life from many aspects. However, human-generated datasets may introduce some human social prejudices to the models (Caliskan-Islam et al., 2016). Recent works have found that ML models can capture, utilize, and even amplify the unintended biases (Zhao et al., 2017), which has raised lots of concerns about the

* Equal contributions from both authors. This work was done when Guanhua Zhang was an intern at Tencent.

Identity-term	Count	Percentage Toxic
gay	868	57.4%
homosexual	202	34.4%
Mexican	116	21.6%
blind	257	14.8%
black	1,123	13.1%
overall	159,686	9.6%

Table 1: Percentage of toxic comments by some specific demographic identity-terms in the dataset released by Dixon et al. (2018).

discrimination problem in NLP models (Sun et al., 2019).

Text classification is one of the fundamental tasks in NLP. It aims at assigning any given sentence to a specific class. In this task, models are expected to make predictions with the semantic information rather than with the demographic group identity information (*e.g.*, “gay”, “black”) contained in the sentences.

However, recent research points out that there widely exist some unintended biases in text classification datasets. For example, in a toxic comment identification dataset released by Dixon et al. (2018), it is found that texts containing some specific identity-terms are more likely to be toxic. More specifically, 57.4% of comments containing “gay” are toxic, while only 9.6% of all samples are toxic, as shown in Table 1.

Because of such a phenomenon, models trained with the dataset may capture the unintended biases and perform differently for texts containing various identity-terms. As a result, predictions of models may discriminate against some demographic minority groups. For instance, sentences like “She makes me happy to be gay” is judged as abusive by models trained on biased datasets in our experiment, which may hinder those minority groups who want to express their feelings on the web freely.

Recent model-agnostic research mitigating the

unintended biases in text classifications can be summarized as data manipulation methods (Sun et al., 2019). For example, Dixon et al. (2018) propose to apply data supplementation with additional labeled sentences to make toxic/non-toxic balanced across different demographic groups. Park et al. (2018) proposes to use data augmentation by applying gender-swapping to sentences with identity-terms to mitigate gender bias. The core of these works is to transform the training sets to an identity-balanced one. However, data manipulation is not always practical. Data supplementation often requires careful selection of the additional sentences *w.r.t.* the identity-terms, the labels, and even the lengths of sentences (Dixon et al., 2018), bringing a high cost for extra data collection and annotation. Data augmentation may result in meaningless sentences (e.g., “He gives birth.”), and is impractical to perform when there are many demographic groups (e.g., for racial bias cases).

In this paper, we propose a model-agnostic debiasing training framework that does not require any extra resources or annotations, apart from a pre-defined set of demographic identity-terms. We tackle this problem from another perspective, in which we treat the unintended bias as a kind of *selection bias* (Heckman, 1979). We assume that there are two distributions, the *non-discrimination distribution*, and the *discrimination distribution* observed in the biased datasets, and every sample of the latter one is drawn independently from the former one following a discrimination rule, *i.e.*, the social prejudice. With such a formalization, mitigating the unintended biases is equivalent to recovering the *non-discrimination distribution* from the selection bias. With a few reasonable assumptions, we prove that we can obtain the unbiased loss of the *non-discrimination distribution* with only the samples from the observed *discrimination distribution* with instance weights. Based on this, we propose a non-discrimination learning framework. Experiments on three datasets show that, despite requiring no extra data, our method is comparable to the data manipulation methods in terms of mitigating the discrimination of models.

The rest of the paper is organized as follows. We summarize the related works in Section 2. Then we give our perspective of the problem and examine the assumptions of commonly-used methods in Section 3. Section 4 introduces our non-discrimination learning framework. Taking three datasets as ex-

amples, we report the experimental results of our methods in Section 5. Finally, we conclude and present the future works in Section 6.

2 Related Works

Non-discrimination and Fairness Non-discrimination focuses on a number of protected demographic groups, and ask for parity of some statistical measures across these groups (Chouldechova, 2017). As mentioned by Friedler et al. (2016), non-discrimination can be achieved only if all groups have similar abilities *w.r.t.* the task in the constructed space which contains the features that we would like to make a decision. There are various kinds of definitions of non-discrimination corresponding to different statistical measures. Popular measures include raw positive classification rate (Calders and Verwer, 2010), false positive and false negative rate (Hardt et al., 2016) and positive predictive value (Chouldechova, 2017), corresponding to different definitions of non-discrimination. Methods like adversarial training (Beutel et al., 2017; Zhang et al., 2018) and fine-tuning (Park et al., 2018) have been applied to remove biasedness.

In the NLP area, fairness and discrimination problems have also gained tremendous attention. Caliskan-Islam et al. (2016) show that semantics derived automatically from language corpora contain human biases. Bolukbasi et al. (2016) show that pre-trained word embeddings trained on large-scale corpus can exhibit gender prejudices and provide a methodology for removing prejudices in embeddings by learning a gender subspace. Zhao et al. (2018) introduce the gender bias problem in coreference resolution and propose a general-purpose method for debiasing.

As for text classification tasks, Dixon et al. (2018) first points out the unintended bias in datasets and proposes to alleviate the bias by supplementing external labeled data. Kiritchenko and Mohammad (2018) examines gender and race bias in 219 automatic sentiment analysis systems and finds that several models show significant bias. Park et al. (2018) focus on the gender bias in abusive language detection task and propose to debias by augmenting the datasets with gender-swapping operation. In this paper, we propose to make models fit a *non-discrimination distribution* with calculated instance weights.

Instance Weighting Instance weighting has been broadly adopted for reducing bias. For example, the Inverse Propensity Score (IPS) (Rosenbaum and Rubin, 1983) method has been successfully applied for causal effect analyses (Austin and Stuart, 2015), selection bias (Schonlau et al., 2009), position bias (Wang et al., 2018; Joachims et al., 2017) and so on. Zadrozny (2004) proposed a methodology for learning and evaluating classifiers under “Missing at Random” (MAR) (Rubin, 1976) selection bias. Zhang et al. (2019) study the selection bias in natural language sentences matching datasets, and propose to fit a leakage-neutral distribution with instance weighting. Jiang and Zhai (2007) propose an instance weighting framework for domain adaptation in NLP, which requires the data of the target domain.

In our work, we formalize the discrimination problem as a kind of “Not Missing at Random” (NMAR) (Rubin, 1976) selection bias from the *non-discrimination distribution* to the *discrimination distribution*, and propose to mitigate the unintended bias with instance weighting.

3 Perspective

In this section, we present our perspective regarding the discrimination problem in text classifications. Firstly, we define what the *non-discrimination distribution* is. Then, we discuss what requirements non-discrimination models should meet and examine some commonly used criteria for non-discrimination. After that, we analyze some commonly used methods for assessing discrimination quantitatively. Finally, we show that the existing debiasing methods can also be seen as trying to recover the *non-discrimination distribution* and examine their assumptions.

3.1 Non-discrimination Distribution

The unintended bias in the datasets is the legacy of the human society where discrimination widely exists. We denote the distribution in the biased datasets as *discrimination distribution* \mathcal{D} .

Given the fact that the real world is discriminatory although it should not be, we assume that there is an ideal world where no discrimination exists, and the real world is merely a biased reflection of the non-discrimination world. Under this perspective, we assume that there is an *non-discrimination distribution* $\hat{\mathcal{D}}$ reflecting the ideal world, and the *discrimination distribution* \mathcal{D} is

drawn from $\hat{\mathcal{D}}$ but following a discriminatory rule, the social prejudice. Attempting to correct the bias of datasets is equivalent to recover the original *non-discrimination distribution* $\hat{\mathcal{D}}$.

For the text classification tasks tackled in this paper, we denote X as the sentences, Y as the (binary) label indicator variable¹, Z as the demographic identity information (e.g. “gay”, “black”, “female”) in every sentence. In the following paper, we use $P(\cdot)$ to represent the probability of the *discrimination distribution* \mathcal{D} in datasets, and $Q(\cdot)$ for *non-discrimination distribution* $\hat{\mathcal{D}}$. Then the non-discrimination distribution $\hat{\mathcal{D}}$ should meet that,

$$Q(Y|Z) = Q(Y),$$

which means that the demographic identity information is independent of the labels².

3.2 Non-Discrimination Model

For text classification tasks, models are expected to make predictions by understanding the semantics of sentences rather than by some single identity-terms. As mentioned in Dixon et al. (2018), a model is defined as biased if it performs better for sentences containing some specific identity-terms than for ones containing others. In other words, a non-discrimination model should perform similarly across sentences containing different demographic groups. However, “perform similarly” is indeed hard to define. Thus, we pay more attention to some criteria defined on demographic groups.

A widely-used criterion is *Equalized Odds* (also known as *Error Rate Balance*) defined by Choudhry (2017), requiring the \hat{Y} to be independent of Z when Y is given, in which \hat{Y} refers to the predictions of the model. This criterion is also used by Borkan et al. (2019) in text classifications.

Besides the *Equalized Odds* criterion, a straightforward criterion for judging non-discrimination is *Statistical Parity* (also known as *Demographic Parity*, *Equal Acceptance Rates*, and *Group Fairness*) (Calders and Verwer, 2010; Dwork et al., 2012), which requires \hat{Y} to be independent of Z ,

¹In this paper, we focus on binary classification problems, but the proposed methodology can be easily extended to multi-class classifications.

²There may be a lot of distributions satisfying the equation. However, as we only focus on the discrimination problem in the text classification task, we suppose that there is a unique *non-discrimination distribution* $\hat{\mathcal{D}}$ which reflects the ideal world in the desired way and the observed biased dataset is drawn from it following a discriminatory rule.

i.e., $\Pr(\hat{Y}|Z) = \Pr(\hat{Y})$. Another criterion is *Predictive Parity* (Chouldechova, 2017), which requires Y to be independent of Z when condition $\hat{Y} = 1$ is given, *i.e.*, $\Pr(Y|\hat{Y} = 1, Z) = \Pr(Y|\hat{Y} = 1)$. Given the definitions of the three criterions, we propose the following theorem, and the proof is presented in Appendix A.

Theorem 1 (Criterion Consistency). *When tested in a distribution in which $\Pr(Y|Z) = \Pr(Y)$, \hat{Y} satisfying Equalized Odds also satisfies Statistical Parity and Predictive Parity.*

Based on the theorem, in this paper, we propose to evaluate models under a distribution where the demographic identity information is not predictive of labels to unify the three widely-used criteria. Specifically, we define that a *non-discrimination model* should meet that,

$$\Pr(\hat{Y}|Y, Z) = \Pr(\hat{Y}|Y),$$

when tested in a distribution where $\Pr(Y|Z) = \Pr(Y)$.

3.3 Assessing the Discrimination

Identity Phrase Templates Test Sets (IPTTS) are widely used as non-discrimination testing sets to assess the models' discrimination (Dixon et al., 2018; Park et al., 2018; Sun et al., 2019; Kiritchenko and Mohammad, 2018). These testing sets are generated by several templates with slots for each of the identity-terms. Identity-terms implying different demographic groups are slotted into the templates, *e.g.*, "I am a boy." and "I am a girl.", and it's easy to find that IPTTS satisfies $\Pr(Y|Z) = \Pr(Y)$. A non-discrimination model is expected to perform similarly in sentences generated by the same template but with different identity-terms.

For metrics, False Positive Equality Difference (FPED) and False Negative Equality Difference (FNED) are used (Dixon et al., 2018; Park et al., 2018), as defined below.

$$\begin{aligned} \text{FPED} &= \sum_z |\text{FPR}_z - \text{FPR}_{\text{overall}}|, \\ \text{FNED} &= \sum_z |\text{FNR}_z - \text{FNR}_{\text{overall}}|, \end{aligned}$$

in which, $\text{FPR}_{\text{overall}}$ and $\text{FNR}_{\text{overall}}$, standing for False Positive Rate and False Negative Rate respectively, are calculated in the whole IPTTS. Correspondingly, FPR_z and FNR_z are calculated on each subset of the data containing each specific

identity-term. These two metrics can be seen as a relaxation of *Equalized Odds* mentioned in Section 3.2 (Borkan et al., 2019).

It should also be emphasized that FPED and FNED do not evaluate the accuracy of models at all, and models can get lower FPED and FNED by making trivial predictions. For example, when tested in a distribution where $\Pr(Y|Z) = \Pr(Y)$, if a model makes the same predictions for all inputs, FPED and FNED will be 0, while the model is completely useless.

3.4 Correcting the Discrimination

Data manipulation has been applied to correct the discrimination in the datasets (Sun et al., 2019). Previous works try to supplement or augment the datasets to an identity-balanced one, which, in our perspective, is primarily trying to recover the *non-discrimination distribution* $\hat{\mathcal{D}}$.

For data supplementation, Dixon et al. (2018) adds some additional non-toxic samples containing those identity-terms which appear disproportionately across labels in the original biased dataset. Although the method is reasonable, due to high cost, it is not always practical to add additional labeled data with specific identity-terms, as careful selection of the additional sentences *w.r.t.* the identity-terms, the labels, and even the lengths of sentences is required (Dixon et al., 2018).

The gender-swapping augmentation is a more common operation to mitigate the unintended bias (Zhao et al., 2018; Sun et al., 2019). For text classification tasks, Park et al. (2018) augment the datasets by swapping the gender-implying identity-terms (*e.g.*, "he" to "she", "actor" to "actress") in the sentences of the training data to remove the correlation between Z and Y . However, it is worth mentioning that the gender-swapping operation additionally assumes that the *non-discrimination distribution* $\hat{\mathcal{D}}$ meets the followings,

$$\begin{aligned} Q(X^\neg|Z) &= Q(X^\neg), \\ Q(Y|X^\neg, Z) &= Q(Y|X^\neg), \end{aligned}$$

in which X^\neg refers to the content of sentences except for the identity information. And we argue that these assumptions may not hold sometimes. For example, the first assumption may result in some meaningless sentences (*e.g.*, "He gives birth.") (Sun et al., 2019). Besides, this method is not practical for situations with many demographic groups.

4 Our Instance Weighting Method

In this section, we introduce the proposed method for mitigating discrimination in text classifications. We first make a few assumptions about how the *discrimination distribution* \mathcal{D} in the datasets are generated from the *non-discrimination distribution* $\hat{\mathcal{D}}$. Then we demonstrate that we can obtain the unbiased loss on $\hat{\mathcal{D}}$ only with the samples from \mathcal{D} , which makes models able to fit the *non-discrimination distribution* $\hat{\mathcal{D}}$ without extra resources or annotations.

4.1 Assumptions about the Generation Process

Considering the perspective that the *discrimination distribution* is generated from the *non-discrimination distribution* $\hat{\mathcal{D}}$, we refer $S \in [0, 1]$ as the selection indicator variable, which indicates whether a sample is selected into the biased dataset or not. Specifically, we assume that every sample (x, z, y, s) ³ is drawn independently from $\hat{\mathcal{D}}$ following the rule that, if $s = 1$ then the sample is selected into the dataset, otherwise it is discarded, then we have

Assumption 1. $P(\cdot) = Q(\cdot|S = 1)$,

and as defined in Section 3.1, the *non-discrimination distribution* $\hat{\mathcal{D}}$ satisfies

Assumption 2. $Q(Y|Z) = Q(Y)$.

Ideally, if the values of S are entirely at random, then the generated dataset can correctly reflect the original *non-discrimination distribution* $\hat{\mathcal{D}}$ and does not have discrimination. However, due to social prejudices, the value of S is not random. Inspired by the fact that some identity-terms are more associated with some specific labels than other identity-terms (e.g., sentences containing ‘‘gay’’ are more likely to be abusive in the dataset as mentioned before), we assume that S is controlled by Y and Z ⁴. We also assume that, given any Z and Y , the conditional probability of $S = 1$ is greater than 0, defined as,

Assumption 3. $Q(S = 1|X, Y, Z) = Q(S = 1|Y, Z) > 0$.

Meanwhile, we assume that the social prejudices will not change the marginal probability distribution of Z , defined as,

³Definitions of x, z and y are in Section 3.1.

⁴As we only focus on the discrimination problem in this work, we ignore selection bias on other variables like *topic* and *domain*.

Assumption 4. $P(Z) = Q(Z)$,

which also means that S is independent with Z in $\hat{\mathcal{D}}$, i.e., $Q(S|Z) = Q(S)$.

Among them, Assumption 1 and 2 come from our problem framing. Assumption 3 helps simplify the problem. Assumption 4 helps establish the *non-discrimination distribution* $\hat{\mathcal{D}}$. Theoretically, when Z is contained in X , which is a common case, consistent learners should be asymptotically immune to this assumption (Fan et al., 2005). A more thorough discussion about Assumption 4 can be found in Appendix B.

4.2 Making Models Fit the Non-discrimination Distribution $\hat{\mathcal{D}}$

Unbiased Expectation of Loss Based on the assumptions above, we prove that we can obtain the loss unbiased to the *non-discrimination distribution* $\hat{\mathcal{D}}$ from the *discrimination distribution* with calculated instance weights.

Theorem 2 (Unbiased Loss Expectation). *For any classifier $f = f(x, z)$, and for any loss function $\Delta = \Delta(f(x, z), y)$, if we use $w = \frac{Q(y)}{P(y|z)}$ as the instance weights, then*

$$E_{x,y,z \sim \mathcal{D}} [w\Delta(f(x, z), y)] = E_{x,y,z \sim \hat{\mathcal{D}}} [\Delta(f(x, z), y)].$$

Then we present the proof for Theorem 2.

Proof. We first present an equation with the weight w , in which we use numbers to denote the assumptions used in each step and *bayes* for the *Bayes’ Theorem*.

$$\begin{aligned} w &= \frac{Q(y)}{P(y|z)} && \left. \vphantom{\frac{Q(y)}{P(y|z)}} \right\} 1 \\ &= \frac{Q(y)}{Q(y|z, S=1)} && \left. \vphantom{\frac{Q(y)}{Q(y|z, S=1)}} \right\} \text{bayes} \\ &= \frac{Q(y)}{Q(S=1|z, y)Q(y|z)/Q(S=1|z)} && \left. \vphantom{\frac{Q(y)}{Q(S=1|z, y)Q(y|z)/Q(S=1|z)}} \right\} 2, 4 \\ &= \frac{Q(S=1)}{Q(S=1|z, y)} && \left. \vphantom{\frac{Q(S=1)}{Q(S=1|z, y)}} \right\} 3, \text{ bayes} \\ &= \frac{Q(S=1)}{Q(x, z, y|S=1)Q(S=1)/Q(x, z, y)} && \left. \vphantom{\frac{Q(S=1)}{Q(x, z, y|S=1)Q(S=1)/Q(x, z, y)}} \right\} 1 \\ &= \frac{Q(x, z, y)}{P(x, z, y)} \end{aligned}$$

Then we have

$$\begin{aligned} &E_{x,z,y \sim \mathcal{D}} [w\Delta(f(x, z), y)] \\ &= \int \frac{Q(x, z, y)}{P(x, z, y)} \Delta(f(x, z), y) dP(x, z, y) \\ &= \int \Delta(f(x, z), y) dQ(x, z, y) \\ &= E_{x,y,z \sim \hat{\mathcal{D}}} [\Delta(f(x, z), y)]. \end{aligned}$$

□

Algorithm 1: Non-discrimination Learning

Input: The dataset $\{x, z, y\}$, the number of fold K for cross prediction and the prior probability $Q(Y = 0)$ and $Q(Y = 1)$

Procedure:

- 01 Train classifiers and use K -fold cross-predictions to estimating $P(y|z)$ with the dataset
 - 02 Calculate the weights $w = \frac{Q(y)}{P(y|z)}$ for all samples
 - 03 Train and validate models using w as the instance weights
-

Non-discrimination Learning Theorem 2 shows that, we can obtain the unbiased loss of the *non-discrimination distribution* $\hat{\mathcal{D}}$ by adding proper instance weights to the samples from the *discrimination distribution* \mathcal{D} . In other words, non-discrimination models can be trained with the instance weights $w = \frac{Q(y)}{P(y|z)}$. As the *discrimination distribution* is directly observable, estimating $P(y|z)$ is not hard. In practice, we can train classifiers and use cross predictions to estimate $P(y|z)$ in the original datasets. Since $Q(y)$ is only a real number indicating the prior probability of $Y \in [0, 1]$ on distribution $\hat{\mathcal{D}}$, we do not specifically make an assumption on it. Intuitively, setting $Q(Y) = P(Y)$ can be a good choice. Considering an non-discrimination dataset where $P(Y|Z) = P(Y)$, the calculated weights $\frac{Q(y)}{P(y|z)}$ should be the same for all samples when we set $Q(Y) = P(Y)$, and thus have little impacts on trained models.

We present the step-by-step procedure for non-discrimination learning in Algorithm 1. Note that the required data is only the biased dataset, and a pre-defined set of demographic identity-terms, with which we can extract $\{x, y, z\}$ for all the samples.

5 Experiments

In this section, we present the experimental results for non-discrimination learning. We demonstrate that our method can effectively mitigate the impacts of unintended discriminatory biases in datasets.

5.1 Dataset Usage

We evaluate our methods on three datasets, including the Sexist Tweets dataset, the Toxicity Comments dataset, and the Jigsaw Toxicity dataset.

Sexist Tweets We use the Sexist Tweets dataset released by Waseem and Hovy (2016); Waseem (2016), which is for abusive language detection

Dataset	Size	Positives	avg. Length
Sexist Tweets	12,097	24.7%	14.7
Toxicity Comments	159,686	9.6%	68.2
Jigsaw Toxicity	1,804,874	8.0%	51.3

Table 2: Statistics of the three datasets for evaluation.

task⁵. The dataset consists of tweets annotated by experts as “sexist” or “normal.” We process the dataset as to how Park et al. (2018) does. It is reported that the dataset has an unintended gender bias so that models trained in this dataset may consider “You are a good woman.” as “sexist.” We randomly split the dataset in a ratio of 8 : 1 : 1 for training-validation-testing and use this dataset to evaluate our method’s effectiveness on mitigating gender discrimination.

Toxicity Comments Another choice is the Toxicity Comments dataset released by Dixon et al. (2018), in which texts are extracted from Wikipedia Talk Pages and labeled by human raters as either toxic or non-toxic. It is found that in this dataset, some demographic identity-terms (*e.g.*, “gay”, “black”) appear disproportionately among labels. As a result, models trained in this dataset can be discriminatory among groups. We adopt the split released by Dixon et al. (2018) and use this dataset to evaluate our method’s effectiveness on mitigating discrimination towards minority groups.

Jigsaw Toxicity We also tested a recently released large-scale dataset Jigsaw Toxicity from Kaggle⁶, in which it is found that some frequently attacked identities are associated with toxicity. Sentences in the dataset are extracted from the Civil Comment platform and annotated with toxicity and identities mentioned in every sentence. We randomly split the dataset into 80% for training, 10% for validation and testing respectively. The dataset is used to evaluate our method’s effectiveness on large-scale datasets.

The statistic of the three datasets is shown as in Table 2.

5.2 Evaluation Scheme

Apart from the original testing set of each dataset, we use the Identity Phrase Templates Test Sets (IPTTS) described in Section 3.3 to evaluate

⁵Unfortunately, due to the rules of Twitter, some TweetIDs got expired, so we cannot collect the exact same dataset as Park et al. (2018).

⁶<https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>

the models as mentioned in Section 3.3. For experiments with the Sexist Tweets dataset, we generate IPTTS following Park et al. (2018). For experiments with Toxicity Comments datasets and Jigsaw Toxicity, we use the IPTTS released by Dixon et al. (2018). Details about the IPTTS generation are introduced in Appendix C.

For metrics, we use FPED and FNED in IPTTS to evaluate how discriminatory the models are, and lower scores indicate better equality. However, as mentioned in Section 3.3, these two metrics are not enough since models can achieve low FPED and FNED by making trivial predictions in IPTTS. So we use AUC in both the original testing set and IPTTS to reflect the trade-off between the debiasing effect and the accuracy of models. We also report the significance test results under confidence levels of 0.05 for Sexist Tweets dataset and Jigsaw Toxicity dataset⁷.

For baselines, we compare with the gender-swapping method proposed by Park et al. (2018) for the Sexist Tweets dataset, as there are only two demographics groups (male and female) provided by the dataset, it’s practical for swapping. For the other two datasets, there are 50 demographics groups, and we compare them with data supplementation proposed by Dixon et al. (2018).

5.3 Experiment Setup

To generate the weights, we use Random Forest Classifiers to estimate $P(y|z)$ following Algorithm 1. We simply set $Q(Y) = P(Y)$ to partial out the influence of the prior probability of Y . The weights are used as the sample weights to the loss functions during training and validation.

For experiments with the Sexist Tweets dataset, we extract the gender identity words (released by Zhao et al. (2018)) in every sentence and used them as Z . For experiments with Toxicity Comments dataset, we take the demographic group identity words (released by Dixon et al. (2018)) contained in every sentence concatenated with the lengths of sentences as Z , just the same as how Dixon et al. (2018) chose the additional sentence for data supplement. For experiments with the Jigsaw Toxicity dataset, the provided identity attributes of every sentence and lengths of sentences are used as Z .

For experiments with the Toxicity Comments dataset, to compare with the results released by

⁷As we use some results from Dixon et al. (2018) directly, we don’t report the significance test results for Toxicity Comments dataset.

Model	Orig. AUC	IPTTS AUC	FPED	FNED
Baseline	0.920	0.673	0.147	0.204
Swap	0.911†	0.651	0.047†	0.050†
Weight	0.897†‡	0.686‡	0.057†	0.086†‡
Baseline ⁺	0.900	0.624	0.049	0.099
Swap ⁺	0.890†	0.611	0.008†	0.013†
Weight ⁺	0.881†‡	0.647†‡	0.007†	0.024†‡

† $p < 0.05$ compared with *Baseline*

‡ $p < 0.05$ compared with *Swap*

Table 3: Experimental results with Sexist Tweets dataset. “+” refers to models using debiased word embeddings.

Model	Orig. AUC	IPTTS AUC	FPED	FNED
Baseline	0.960	0.952	7.413	3.673
Supplement	0.959	0.960	5.294	3.073
Weight	0.956	0.961	4.798	2.491

The results of *Baseline* and *Supplement* are taken from Dixon et al. (2018)

Table 4: Experimental results with Toxicity Comments dataset.

Model	Orig. AUC	IPTTS AUC	FPED	FNED
Baseline	0.928	0.993	3.088	3.317
Supplement	0.928	0.999†	0.180†	3.111
Weight	0.922†‡	0.999†	0.085†	2.538

† $p < 0.05$ compared with *Baseline*

‡ $p < 0.05$ compared with *Supplement*

Table 5: Experimental results with Jigsaw Toxicity dataset.

Dixon et al. (2018), we use their released codes, where a three-layer Convolutional Neural Network (CNN) model is used. For experiments with Sexist Tweets dataset and Jigsaw Toxicity dataset, as our method is model-agnostic, we simply implement a one-layer LSTM with a dimensionality of 128 using Keras and Tensorflow backend.⁸

For all models, pre-trained GloVe word embeddings (Pennington et al., 2014) are used. We also report results when using gender-debiased pre-trained embeddings (Bolukbasi et al., 2016) for experiments with Sexist Tweets. All the reported results are the average numbers of ten runs with different random initializations.

5.4 Experimental Results

In this section, we present and discuss the experimental results. As expected, training with calculated weights can effectively mitigate the impacts of the unintended bias in the datasets.

Sexist Tweets Tabel 3 reports the results on Sexist Tweets dataset. *Baseline* refers to vanilla mod-

⁸Codes are publicly available at <https://github.com/ghzhang233/Non-Discrimination-Learning-for-Text-Classification>.

els. *Swap* refers to models trained and validated with 2723 additional gender-swapped samples to balance the identity-terms across labels (Park et al., 2018). *Weight* refers to models trained and validated with calculated weights. “+” refers to models using debiased word embeddings.

Regarding the results with the GloVe word embeddings, we can find that *Weight* performs significantly better than *Baseline* under FPED and FNED, which demonstrate that our method can effectively mitigate the discrimination of models. *Swap* outperforms *Weight* in FPED and FNED, but our method achieves significantly higher IPTTS AUC. We notice that *Swap* even performs worse in terms of IPTTS AUC than *Baseline* (although the difference is not significant at 0.05), which implies that cost for the debiasing effect of *Swap* is the loss of models’ accuracy, and this can be ascribed to the gender-swapping assumptions as mentioned in Section 3.4. We also notice that both *Weight* and *Swap* have lower Orig. AUC than *Baseline* and this can be ascribed to that the unintended bias pattern is mitigated.

Regarding the results with the debiased word embeddings, the conclusions remain largely unchanged, while *Weight* get a significant improvement over *Baseline* in terms of IPTTS AUC. Besides, compared with GloVe embeddings, we can find that debiased embeddings can effectively improve FPED and FNED, but Orig. AUC and IPTTS AUC also drop.

Toxicity Comments Table 4 reports the results on Toxicity Comments dataset. *Baseline* refers to vanilla models. *Supplement* refers to models trained and validated with 4620 additional samples to balance the identity-terms across labels (Dixon et al., 2018). *Weight* refers to models trained and validated with calculated instance weights.

From the table, we can find that *Weight* outperforms *Baseline* in terms of IPTTS AUC, FPED, and FNED, and also gives slightly better debiasing performance compared with *Supplement*, which demonstrate that the calculated weights can effectively make models more non-discriminatory. Meanwhile, *Weight* performs similarly in Orig. AUC to all the other methods, indicating that our method does not hurt models’ generalization ability very much.

In general, the results demonstrate that our method can provide a better debiasing effect without additional data, and avoiding the high cost of

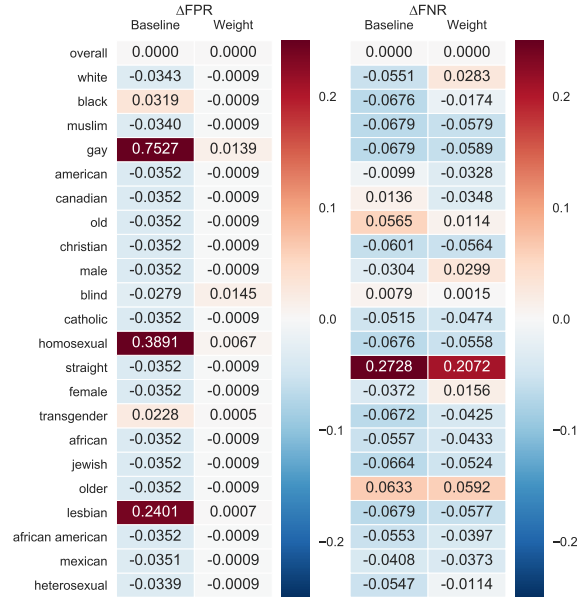


Figure 1: Comparison for the evaluation results of *Baseline* and *Weight* for sentences containing a selection of specific identities in IPTTS in Jigsaw Toxicity dataset, in which $\Delta FPR_z = FPR_z - FPR_{overall}$, and $\Delta FNR_z = FNR_z - FNR_{overall}$. Values closer to 0 indicate better equality. Best viewed in color.

extra data collection and annotation makes it more practical for adoptions.

Jigsaw Toxicity Table 5 reports the results on Jigsaw Toxicity dataset. *Baseline* refers to vanilla models. *Supplement* refers to models trained and validated with 15249 additional samples extracted from Toxicity Comments to balance the identity-terms across labels. *Weight* refers to models trained with calculated weights.

Similar to results on Toxicity Comments, we find that both *Weight* and *Supplement* perform significantly better than *Baseline* in terms of IPTTS AUC and FPED, and the results of *Weight* and *Supplement* are comparable. On the other hand, we notice that *Weight* and *Supplement* improve FNED slightly, while the differences are not statistically significant at confidence level 0.05.

To gain better knowledge about the debiasing effects, we further visualize the evaluation results on the Jigsaw Toxic dataset for sentences containing some specific identity-terms in IPTTS in Figure 1, where ΔFPR_z and ΔFNR_z are presented. Based on the definition of FPED and FNED, values closer to 0 indicate better equality. We can find that *Baseline* trained in the original biased dataset can discriminate against some demographic groups.

For example, sentences containing identity words like “gay”, “homosexual” and “lesbian” are more likely to be falsely judged as “toxic” as indicated by ΔFPR , while ones with words like “straight” are more likely to be falsely judged as “not toxic” as indicated by ΔFNR . We can also notice that *Weight* performs more consistently among most identities in both FPR and FNR. For instance, ΔFPR of the debiased model in samples with “gay”, “homosexual” and “lesbian” significantly come closer to 0, while $|\Delta\text{FNR}|$ also drop for “old” and “straight”.

We also note that $\text{FPR}_{\text{overall}}$ and $\text{FNR}_{\text{overall}}$ of *Weight* are significantly better than the results of *Baseline*, i.e., $\text{FPR}_{\text{overall}}$ results are 0.001 and 0.068 for *Weight* and *Baseline* respectively, and $\text{FNR}_{\text{overall}}$ results are 0.061 and 0.068 for *Weight* and *Baseline* respectively, representing that *Weight* is both more accurate and more non-discriminatory on the IPTTS set.

6 Conclusion

In this paper, we focus on the unintended discrimination bias in existing text classification datasets. We formalize the problem as a kind of selection bias from the *non-discrimination distribution* to the *discrimination distribution* and propose a debiasing training framework that does not require any extra resources or annotations. Experiments show that our method can effectively alleviate discrimination. It’s worth mentioning that our method is general enough to be applied to other tasks, as the key idea is to obtain the loss on the *non-discrimination distribution*, and we leave this to future works.

Acknowledgments

Conghui Zhu and Tiejun Zhao are supported by National Key R&D Program of China (Project No. 2017YFB1002102).

References

Peter C Austin and Elizabeth A Stuart. 2015. Moving towards best practice when using inverse probability of treatment weighting (iptw) using the propensity score to estimate causal treatment effects in observational studies. *Statistics in medicine*, 34(28):3661–3679.

Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. 2007. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144.

Alex Beutel, Ed H Chi, Jilin Chen, and Zhe Zhao. 2017. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*.

Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*, pages 4349–4357.

Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 491–500. ACM.

Toon Calders and Sicco Verwer. 2010. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292.

Aylin Caliskan-Islam, Joanna J. Bryson, and Arvind Narayanan. 2016. Semantics derived automatically from language corpora necessarily contain human biases. *Science*, 356(6334):183–186.

Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163.

Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 67–73. ACM.

Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226. ACM.

Wei Fan, Ian Davidson, Bianca Zadrozny, and Philip S Yu. 2005. An improved categorization of classifier’s sensitivity on sample selection bias. In *Fifth IEEE International Conference on Data Mining (ICDM’05)*, pages 4–pp. IEEE.

Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236*.

Moritz Hardt, Eric Price, Nati Srebro, et al. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pages 3315–3323.

James J Heckman. 1979. Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pages 153–161.

- Jing Jiang and ChengXiang Zhai. 2007. Instance weighting for domain adaptation in nlp. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pages 264–271.
- Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 781–789. ACM.
- Svetlana Kiritchenko and Saif Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53.
- Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. Reducing gender bias in abusive language detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2799–2804.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Donald B Rubin. 1976. Inference and missing data. *Biometrika*, 63(3):581–592.
- Matthias Schonlau, Arthur Van Soest, Arie Kapteyn, and Mick Couper. 2009. Selection bias in web surveys and the use of propensity scores. *Sociological Methods & Research*, 37(3):291–318.
- Hidetoshi Shimodaira. 2000. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating gender bias in natural language processing: Literature review. In *Proceedings of the 57th Annual Meeting of the Association of Computational Linguistics*, pages 1630–1640.
- Xuanhui Wang, Nadav Golbandi, Michael Bendersky, Donald Metzler, and Marc Najork. 2018. Position bias estimation for unbiased learning to rank in personal search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 610–618. ACM.
- Zeeraak Waseem. 2016. Are you a racist or am i seeing things? annotator influence on hate speech detection on twitter. In *Proceedings of the first workshop on NLP and computational social science*, pages 138–142.
- Zeeraak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93.
- Bianca Zadrozny. 2004. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, page 114. ACM.
- Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340. ACM.
- Guanhua Zhang, Bing Bai, Jian Liang, Kun Bai, Shiyu Chang, Mo Yu, Conghui Zhu, and Tiejun Zhao. 2019. Selection bias explorations and debias methods for natural language sentence matching datasets. In *Proceedings of the 57th Annual Meeting of the Association of Computational Linguistics*, pages 4418–4429.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 2.

A Proof for the Criterion Consistency Theorem

Proof. Here we present the proof for Theorem 1. For the Statistical Parity criterion,

$$\begin{aligned}
 \Pr(\hat{Y}|Z) &= \sum_{y \in [0,1]} \Pr(\hat{Y}, Y = y|Z) \\
 &= \sum_{y \in [0,1]} \Pr(\hat{Y}|Y = y, Z) \Pr(Y = y|Z) \\
 &= \sum_{y \in [0,1]} \Pr(\hat{Y}|Y = y) \Pr(Y = y) \\
 &= \Pr(\hat{Y}).
 \end{aligned}$$

For the Predictive Parity criterion,

$$\begin{aligned}
 \Pr(Y|\hat{Y} = 1, Z) &= \frac{\Pr(\hat{Y} = 1|Y, Z) \Pr(Y|Z)}{\Pr(\hat{Y} = 1|Z)} \\
 &= \frac{\Pr(\hat{Y} = 1|Y) \Pr(Y)}{\Pr(\hat{Y} = 1)} \\
 &= \Pr(Y|\hat{Y} = 1).
 \end{aligned}$$

□

B Discussion about Assumption 4

We show that even if the assumption does not hold, we can still make models fit $Q(Y|X)$ with calculated weights when Z is contained in X , which is the common setting in practical.

We firstly present the equation of the weights w without the assumption $P(Z) = Q(Z)$.

$$\begin{aligned}
 w &= \frac{Q(y)}{P(y|z)} && \left. \begin{array}{l} \\ \\ \end{array} \right\} 1 \\
 &= \frac{Q(y)}{Q(y|z, S=1)} && \left. \begin{array}{l} \\ \\ \end{array} \right\} \text{bayes} \\
 &= \frac{Q(y)}{Q(S=1|z, y)Q(y|z)/Q(S=1|z)} && \left. \begin{array}{l} \\ \\ \end{array} \right\} 2, 3 \\
 &= \frac{Q(S=1|z)}{Q(S=1|x, z, y)} && \left. \begin{array}{l} \\ \\ \end{array} \right\} \text{bayes} \\
 &= \frac{Q(S=1)}{Q(S=1|x, z, y)} \cdot \frac{Q(z|S=1)}{Q(z)} && \left. \begin{array}{l} \\ \\ \end{array} \right\} 1, \text{bayes} \\
 &= \frac{Q(x, z, y)}{P(x, z, y)} \cdot \frac{P(z)}{Q(z)}
 \end{aligned}$$

After applying these weights to every sample in the dataset, we can get a new distribution defined as below,

$$P^*(x, y, z_x) = \frac{w_{x, y, z_x} \cdot P(x, y, z_x)}{\int w_{x', y', z'_x} \cdot dP(x', y', z'_x)}$$

in which we use $P^*(\cdot)$ to represent the probability in the obtained distribution. As Z is contained in X , we use Z_X to represent the specific Z contained in every X .

Then we have

$$\begin{aligned}
 P^*(y|x) &= \frac{P^*(x, z_x, y)}{\sum_{y'} P^*(x, z_x, y')} \\
 &= \frac{P(x, z_x, y) \cdot \frac{Q(x, z_x, y)}{P(x, z_x, y)} \cdot \frac{P(z_x)}{Q(z_x)}}{\sum_{y'} P(x, z_x, y') \cdot \frac{Q(x, z_x, y')}{P(x, z_x, y')} \cdot \frac{P(z_x)}{Q(z_x)}} \\
 &= \frac{\frac{P(z_x)}{Q(z_x)} \cdot Q(x, z_x, y)}{\frac{P(z_x)}{Q(z_x)} \cdot \sum_{y'} Q(x, z_x, y')} \\
 &= Q(y|x, z_x) \\
 &= Q(y|x),
 \end{aligned}$$

and

$$\begin{aligned}
 P^*(x) &= \sum_y P^*(x, z_x, y) \\
 &= \sum_y \frac{w_{x, z_x, y} \cdot P(x, z_x, y)}{\int w_{x', z_{x'}, y'} \cdot dP(x', z_{x'}, y')} \\
 &\propto \sum_y \frac{Q(x, z_x, y)}{P(x, z_x, y)} \cdot \frac{P(z_x)}{Q(z_x)} \cdot P(x, z_x, y) \\
 &= \frac{P(z_x)}{Q(z_x)} \cdot Q(x, z_x) \\
 &= \frac{P(z_x)}{Q(z_x)} \cdot Q(x).
 \end{aligned}$$

Given the result $P^*(y|x) = Q(y|x)$, the consistent learners should be asymptotically immune to different assumptions regarding $Q(Z)$, where a

learner is defined as consistent if the learning algorithm can find a model θ that is equivalent to the true model at producing class conditional probabilities given an exhaustive training data set (Fan et al., 2005). In practical, however, as the requirements are often hard to met, we note that models may still be affected due to the deviation between $P^*(x)$ and $Q(x)$, which is widely studied in the covariate shift problem (Shimodaira, 2000; Ben-David et al., 2007; Jiang and Zhai, 2007). In our paper, as we don't assume the availability of extra resources and prior knowledge, we simply set $P(Z) = Q(Z)$. We leave more explorations about this assumption for future work.

C Details about the IPTTS Generation

Templates	Label
(verb inoffensive) (identity-term).	0
(verb offensive) (identity-term).	1
I am (identity-term).	0
I hate (identity-term).	1
Being (identity-term) is (adj. inoffensive).	0
Being (identity-term) is (adj. offensive).	1
You are a (adj. inoffensive) (identity-term).	0
You are a (adj. offensive) (identity-term).	1

Table 6: Templates used to generate IPTTS.

Type	Example Words
adj. inoffensive	great, fun, nice, neat
adj. offensive	disgusting, filthy, nasty, rotten
verb inoffensive	hug, love, like, respect
verb offensive	kill, murder, hate, destroy
male identity	actor, airman, boy, man
female identity	actress, airwoman, girl, woman

Table 7: Examples of slotted words to generate IPTTS.

For experiments with the Sexist Tweets dataset, we generate IPTTS following Park et al. (2018). The templates used are the same as Park et al. (2018), as shown in Table 6. We use the released codes by Dixon et al. (2018) and use the gender word pairs released by Zhao et al. (2018) as "identity-term." Some of the slotted words are presented in Table 7. To make sentences longer, we also add some semantic-neutral sentences provided by Dixon et al. (2018) as a suffix to each template. Finally, we get 75238 samples, 37538 of which are abusive, and the mean of sentence lengths is 17.5.

For experiments with Toxicity Comments datasets and Jigsaw Toxicity, we use the IPTTS

Term	Origin			Weight		
	Toxic	Overall	Δ	Toxic	Overall	Δ
white	5.98	2.13	3.85	2.89	2.14	0.75
black	3.10	1.07	2.03	1.22	1.07	0.15
muslim	1.57	0.58	0.99	0.58	0.59	-0.01
gay	1.29	0.35	0.94	0.39	0.34	0.05
american	2.70	2.11	0.59	2.76	2.13	0.63
canadian	1.38	1.82	-0.44	1.48	1.82	-0.34
old	2.62	2.18	0.44	2.63	2.18	0.45
christian	0.89	0.54	0.35	0.73	0.55	0.18
male	0.73	0.44	0.29	0.41	0.45	-0.04
blind	0.51	0.28	0.23	0.55	0.28	0.27
catholic	0.63	0.82	-0.19	0.65	0.83	-0.18
homosexual	0.26	0.08	0.18	0.09	0.08	0.01
straight	0.51	0.37	0.14	0.46	0.38	0.08
female	0.50	0.37	0.13	0.28	0.37	-0.09
transgender	0.21	0.09	0.12	0.10	0.09	0.01
african	0.30	0.20	0.10	0.20	0.20	0.00
jewish	0.28	0.19	0.09	0.17	0.19	-0.02
older	0.16	0.25	-0.09	0.15	0.25	-0.10
lesbian	0.11	0.03	0.08	0.03	0.03	0.00
african american	0.16	0.09	0.07	0.09	0.10	-0.01
mexican	0.20	0.13	0.07	0.17	0.13	0.04
heterosexual	0.09	0.03	0.06	0.03	0.03	0.00

Table 8: Frequency of a selection of identity-terms in toxic samples and overall in Jigsaw Toxicity dataset. % is omitted.

released by [Dixon et al. \(2018\)](#). The testing set is created by several templates slotted by a broad range of identity-terms, which consists of 77000 examples, 50% of which are toxic.

D Frequency of Identity-terms in Toxic Samples and Overall

To give a better understanding of how the weights change the distribution of datasets, we compare the original Jigsaw Toxicity dataset and the one with calculated weights for the frequency of a selection of identity-terms in toxic samples and overall, as shown in Table 8.

We can find that after adding weights, the gap between frequency in toxic samples and overall significantly decrease for almost all identity-terms, which demonstrate that the unintended bias in datasets is effectively mitigated.