

The Authors Matter: Understanding and Mitigating Implicit Bias in Deep Text Classification

Haochen Liu¹, Wei Jin¹, Hamid Karimi¹, Zitao Liu^{2*}, Jiliang Tang¹

¹ Michigan State University, East Lansing, MI, USA

² TAL Education Group, Beijing, China

{liuhaoc1,jinwei2,karimiha}@msu.edu; liuzitao@100tal.com; tangjili@msu.edu

Abstract

It is evident that deep text classification models trained on human data could be biased. In particular, they produce biased outcomes for texts that explicitly include identity terms of certain demographic groups. We refer to this type of bias as explicit bias, which has been extensively studied. However, deep text classification models can also produce biased outcomes for texts written by authors of certain demographic groups. We refer to such bias as implicit bias, of which we still have a rather limited understanding. In this paper, we first demonstrate that implicit bias exists in different text classification tasks for different demographic groups. Then, we build a learning-based interpretation method to deepen our knowledge of implicit bias. Specifically, we verify that classifiers learn to make predictions based on language features that are related to the demographic attributes of the authors. Next, we propose a framework **Debiased-TC** to train deep text classifiers to make predictions on the right features and consequently mitigate implicit bias. We conduct extensive experiments on three real-world datasets. The results show that the text classification models trained under our proposed framework outperform traditional models significantly in terms of fairness, and also slightly in terms of classification performance.

1 Introduction

Many recent studies have suggested that machine learning algorithms can learn social prejudices from data produced by humans, and thereby show systemic bias in performance towards specific demographic groups or individuals (Mehrabi et al., 2019; Blodgett et al., 2020; Shah et al., 2020). As one machine learning application, text classification has been proven to be discriminatory towards certain groups of people (Dixon et al., 2018;

Borkan et al., 2019). Text classification applications such as sentiment analysis and hate speech detection are common and widely used in our daily lives. If a biased hate speech detection model is deployed by a social media service provider to filter users' comments, the comments related to different demographic groups can have uneven chances to be recognized and removed. Such a case will cause unfairness and bring in negative experiences to users. Thus, it is highly desired to mitigate the bias in text classification.

The majority of existing studies on bias and fairness in text classification have mainly focused on the bias towards the individuals mentioned in the text content. For example, in (Dixon et al., 2018; Park et al., 2018; Zhang et al., 2020), it is investigated how text classification models perform unfairly on texts containing demographic identity terms such as "gay" and "muslim". In such scenarios, the demographic attributes of the individuals subject to bias explicitly exist in the text. In this work, we refer to this kind of bias as **explicit bias**. Bias in texts, however, can be reflected more subtly and insidiously. While a text may not contain any reference to a specific group or individual, the content can somehow be revealing of the demographic information of the author. As shown in (Coulmas, 2013; Preoțiuc-Pietro and Ungar, 2018), the language style (e.g., wordings and tone) of a text can be highly correlated with its author's demographic attributes (e.g., age, gender, and race). We find that a text classifier can learn to associate the content with demographic information and consequently make unfair decisions towards certain groups. We refer to such bias as **implicit bias**. Table 1 demonstrates an example of implicit bias. There are two short texts where the first text is written by a white American and the second one is written by an African American. The task is to predict the sentiment of a text by a convolutional

* The corresponding author: Zitao Liu

Table 1: An illustrative example on the implicit bias of a CNN text classification model.

Author	Text	Label	Prediction
White American	Can't wait to visit your new home. Yes, I going to be a great guest!	positive	positive
African American	Can't wait to visit your new home. Yup, I goin to be a great guest!	positive	negative

neural network (CNN) model. Words with a red background indicate those with the salient predictive capability by the model where the darker the color, the more salient the words. The words “yup” and “goin” in the second text are commonly used by African Americans (Liu et al., 2020a) and are irrelevant to the sentiment. However, the CNN model has hinted at them and consequently has predicted a positive text to be negative.

In this work, we aim to understand and mitigate implicit bias in deep text classification models. One key source of bias is the imbalance of training data (Dixon et al., 2018; Park et al., 2018). Thus, existing debiasing methods mainly focus on balancing the training data, such as adding new training data (Dixon et al., 2018) and augmenting data based on identity-term swap (Park et al., 2018). However, these methods cannot be directly applied to mitigate implicit bias. Obtaining new texts from authors of various demographic groups is very expensive. It requires heavy human labor. Meanwhile, given that there is no explicit demographic information in texts, identity-term swap data augmentation is not applicable. Thus, we propose to enhance deep text classification models to mitigate implicit bias in the training process. To achieve this goal, we face tremendous challenges. First, to mitigate the implicit bias, we have to understand how deep models behave. For example, how they correlate implicit features in text with demographic attributes and how the models make biased predictions. Second, we need to design new mechanisms to take advantage of our understandings to mitigate the implicit bias in deep text classifiers.

To address the above challenges, in this paper, we first propose an interpretation method, which sheds light on the formation mechanism of implicit bias in deep text classification models. We show that the implicit bias is caused by the fact that the models make predictions based on incorrect language features in texts. Second, based on this finding, we propose a novel framework **Debiased-TC** (Debiased Text Classification) to mitigate the implicit bias of deep text classifiers. More specifically, we equip the deep classifiers with an additional saliency selection layer that first determines the

correct language features which the model should base on to make predictions. We also propose an optimization method to train the classifiers with the saliency selection layer. Note that both our proposed interpretation method and the learning framework are model-agnostic, which means that they can be applied to any deep text classifier. We evaluate the framework with two popular deep text classification models across various text classification tasks on three public datasets. The experimental results demonstrate that our method significantly mitigates the implicit bias in the classification models while maintaining or even improving their prediction performance.

2 Preliminary Study

In this section, we perform a preliminary study to validate the existence of implicit bias in deep text classification models. We first introduce the data and text classification tasks, and then present the empirical results.

2.1 Data and Tasks

In the preliminary study, we investigate different text classification tasks and various demographic groups to validate the implicit bias. We use three datasets, including the DIAL and PAN16 datasets processed by (Elazar and Goldberg, 2018) and the Multilingual Twitter Corpus (MTC) introduced in (Huang et al., 2020).

The DIAL dataset contains dialectal texts collected from Twitter. Each tweet’s text is associated with the *race* of the author as the demographic attribute, denoted as “white” or “black”, respectively. This dataset is annotated for two classification tasks: sentiment analysis and mention detection. The sentiment analysis task aims to categorize a text as “happy” or “sad”. The mention detection task tries to determine whether a tweet mentions another user, which can also be viewed as distinguishing conversational tweets from non-conversational ones.

The PAN16 dataset consists of tweets. For each tweet, *age* and *gender* of its author have been manually labelled. The demographic attribute age has two categories of “18-34” and “ ≥ 35 ”, and gender

has “male” and “female”. Also, this dataset is annotated for the mention detection task as described above.

The MTC dataset contains multilingual tweets for the hate speech detection task. Each tweet is annotated as “hate speech” or “non hate speech” and associated with four author’s demographic attributes: race, gender, age, and country. We only use the English corpus with the attribute *race*. In this dataset, the attribute race has two categories, i.e., “white” and “nonwhite”.

More information about these three datasets, including their statistical information, annotation process, and the links to downloadable versions of the data can be found in Appendix A.

2.2 Empirical study

In this subsection, we aim to empirically study if text classification models make the predictions dependent on the demographic attributes of the authors of the texts. The explicit bias in text classification tasks stems from the imbalance of training data (Dixon et al., 2018; Park et al., 2018). For example, when there are more negative examples from one group in the training data, the model learns to correlate that group with the negative label, which results in bias. Inspired by this observation, to validate the existence of implicit bias, we investigate if the imbalance of training data in terms of demographic attributes of the authors can lead to biased predictions. To answer this question, we consider the following setting: (1) the training data has an equal number of positive and negative examples; and (2) positive and negative examples in the training data are imbalanced among different groups of the authors according to their demographic attributes. Intuitively, if the predictions are independent of the demographic attributes of authors, the model should still perform similarly for different groups.

For each task and demographic attribute of authors, we consider two labels (i.e., positive and negative) and two demographic groups (i.e., Group I and Group II). For each dataset, we follow the aforementioned setting to build a training set. We make the training set overall balanced in terms of the labels and demographic groups. That is, we set the overall ratio of positive and negative examples as 1:1, and the overall ratio of examples from Group I and Group II as 1:1 as well. Meanwhile, we make the data in each group imbalanced. In

particular, for Group I, we set the ratio of its positive and negative examples to 4:1, while the ratio is automatically set to 1:4 for Group II. We name the proportion of positive and negative samples in Group I as the “balance rate”. We train a CNN text classifier as a representative model on the training set and evaluate it on the test set. We use the false positive/negative rates (Dixon et al., 2018) and the demographic parity rate (a.k.a., positive outcome rate, the probability of the model predicting a positive outcome for one group) (Dwork et al., 2012; Kusner et al., 2017) to evaluate the fairness of the classification models.

The results are shown in Table 2. For the demographic attribute race, Group I/Group II stands for white/black in the DIAL dataset, and white/nonwhite in the MTC dataset. For gender and age, Group I/Group II stands for male/female and age ranges (18-34)/(≥ 35), respectively. From the table, we observe that in terms of different tasks and demographic attributes of authors, the model shows significant bias with the same pattern. For all cases, the demographic group with more positive examples (Group I) always gets a higher false positive rate, a lower false negative rate, and a higher demographic parity rate than the other group. This demonstrates that imbalanced data can cause implicit bias, and the predictions are not independent of the demographic attributes of authors. Since the text itself doesn’t explicitly contain any demographic information, the model could learn to recognize the demographic attributes of authors based on implicit features such as language styles and associate them with a biased outcome. Next, we will understand one formation of implicit bias and then propose Debaised-TC to mitigate it.

3 Understanding Implicit Bias

In this section, we aim to understand the possible underlying formation mechanism of implicit bias. Our intuition is – when a training set for sentiment analysis has more positive examples from white authors and more negative examples from black authors, a classification model trained on such a dataset may learn a “shortcut” (Mahabadi et al., 2020) to indiscriminately associates the language style features of white people with the positive sentiment and those of black people with the negative sentiment. In other words, the model does not use the correct language features (e.g., emotional words) to make the prediction. Thus, we attempt

Table 2: Preliminary study.

Dataset	Task	Demo	False Positive (%)		False Negative (%)		Demographic Parity (%)	
			Group I	Group II	Group I	Group II	Group I	Group II
DIAL	Sentiment	Race	46.97	23.38	21.29	62.75	62.84	30.32
	Mention	Race	48.72	15.99	17.32	34.90	65.70	40.55
PAN16	Mention	Gender	23.90	12.30	13.06	23.01	55.42	44.64
	Mention	Age	24.91	9.88	16.48	26.43	54.22	41.72
MTC	Hate Speech	Race	80.33	1.77	12.13	49.35	84.10	26.21

to examine the following hypothesis: *A deep text classification model presents implicit bias since it makes predictions based on language features that should be irrelevant to the classification task but are correlated with a certain demographic group of authors.* To verify this hypothesis, we first propose an interpretation method to detect the salient words a text classification model relies on to make the prediction. The interpretation model enables us to check the overlapping between the salient words and the words related to the authors’ demographic attributes. Consequently, it allows us to understand the relationship between such overlapping and the model’s implicit bias.

3.1 An Interpretation Method

We follow the idea of the learning-based interpretation method L2X (Chen et al., 2018) to train an explainer to interpret a given model. The reasons for choosing L2X are – 1) as a learning-based explainer, it learns to globally explain the behavior of a model, instead of explaining a single instance at one time; and 2) the explainer has the potential to be integrated into our debiasing framework to mitigate implicit bias in an end-to-end manner, which will be introduced in Section 4.

A binary text classification model $\mathcal{M} : X \rightarrow Y$ maps an input text $X = (x_1, x_2, \dots, x_n)$ to a label $Y \in \{0, 1\}$. For a certain model \mathcal{M} , we seek to specify the contribution of each word in X for \mathcal{M} to make the prediction Y . The contributions can be denoted as a saliency distribution $S = (s_1, s_2, \dots, s_n)$, where s_i is the saliency score of the word x_i , and $\sum_{i=1}^n s_i = 1$. Given a model \mathcal{M} , we train an explainer $\mathcal{E}^{\mathcal{M}} : X \rightarrow S$ to estimate the saliency distribution S of an input text X .

The explainer is trained by maximizing $I(X_S, Y)$, the mutual information (Cover, 1999) between the response variable Y and the selected feature X_S of X under saliency distribution S . The selected feature $X_S = X \odot S = (s_1 \cdot x_1, s_2 \cdot x_2, \dots, s_n \cdot x_n)$ ¹ is calculated as the element-wise

¹Without confusion, we use x_i to denote both a word and its word embedding vector.




Figure 1: An illustration of the bias interpretation model.

product between X and S . In our implementation, we parametrize the explainer by a bi-directional recurrent neural network (RNN) followed by a linear layer and a Softmax layer. More details about the optimization of the explainer can be found in Appendix B.

3.2 Saliency Correlation Measurement

In this work, we assume that the text classification task is totally independent of the demographic attribute of the author of the text. In other words, language features that reflect the author’s demographic information should not be taken as evidence for the main task. Thus, we propose to understand the implicit bias of a deep text classification model by examining the overlapping between salient words for the main task and the words correlated with the demographic attribute.

With the interpretation model, we can estimate the saliency distributions of the input words for the classification task and the demographic attribute prediction task, respectively, and then check their overlapping. As shown in Figure 1, we train two models \mathcal{M}^Y and \mathcal{M}^Z with the same architecture for the former and the latter tasks, respectively. Then, two corresponding explainers \mathcal{E}^Y and \mathcal{E}^Z are trained for them. Thus, given an input text X , two explainers can estimate the saliency distributions S^Y and S^Z on two tasks, respectively. We use the Jensen-Shannon (JS) divergence $JS(S^Y || S^Z)$ to measure the overlap between language features that

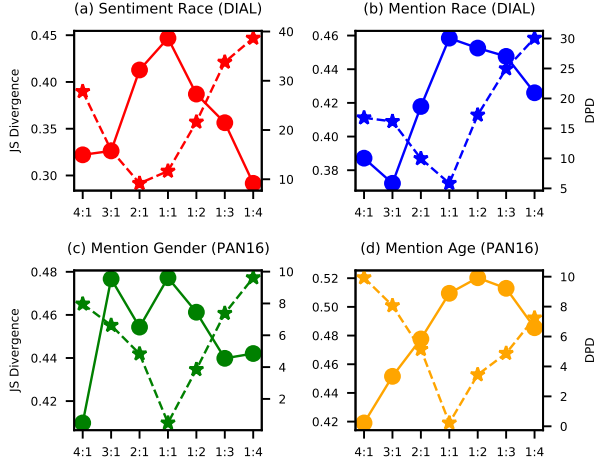


Figure 2: The average JS divergence (solid lines) and DPD (dash lines) vs. the balance rate. The x-axis indicates the balance rate of the training set. The y-axis on the left hand indicates the average JS divergence, and the y-axis on the right hand is the DPD.

these two tasks relying on to make the predictions on Y and Z .

3.3 Empirical Analysis

In this subsection, we present the experiments to verify our hypothesis on the formulation of implicit bias. Following the experimental settings in Section 2.2, we vary the “balance rate” of the training data and then observe how the saliency correlation changes. We use CNN text classifiers (see Appendix C for details) for both \mathcal{M}^Y and \mathcal{M}^Z . In Figure 2, we show how the average JS divergence and the demographic parity difference (DPD) vary with the changes of the balance rate. DPD is the absolute value of the difference between the demographic parity rates for the two groups. We only report the results for DIAL and PAN16 datasets and DPD as the fairness metric since we achieved similar results for other settings. For each task and each demographic attribute, the DPD is small when the training data are balanced and becomes large when the data are imbalanced. However, the JS divergence is large for balanced data while small for imbalanced data. A larger DPD indicates stronger implicit bias and a smaller JS divergence stands for a stronger overlap between the saliency distributions for the two tasks. Thus, these observations suggest that when the training data are imbalanced, the text classifiers tend to use language features related to the demographic attribute of authors to make the prediction.




Figure 3: An illustration of the bias mitigation model.

4 The Bias Mitigation Framework

In the previous section, we showed that a model with implicit bias tends to utilize features related to the demographic attribute of authors to make the prediction, especially when training data is imbalanced in terms of the demographic attribute of authors. One potential solution is to balance the training data by augmenting more examples from underrepresented groups. However, collecting new data from authors of different demographics is expensive. Thus, to mitigate the implicit bias, we propose a novel framework **Debiased-TC**. Our proposed approach can mitigate implicit bias by automatically correcting their selection of input features. In this section, we will introduce the proposed framework with the corresponding optimization method.

4.1 Debiased Text Classification Model

An illustration of Debiased-TC is shown in Figure 3. Similar to the explainer in the interpretation model, we equip the base model \mathcal{M}^Y with a corrector layer \mathcal{C} after the input layer. The corrector $\mathcal{C} : X \rightarrow S$ learns to correct the model’s feature selection. It first maps an input text $X = (x_1, x_2, \dots, x_n)$ to a saliency distribution $S = (s_1, s_2, \dots, s_n)$, which is expected to give high scores to words related to the main tasks and low scores to words related to demographic attributes of authors. Then, it assigns weights to the input features with the saliency scores by calculating $X_S = X \odot S$, which is fed into the classification model \mathcal{M}^Y for prediction.

To train a corrector to achieve the expected goal, we adopt the idea of adversarial training. More specifically, in addition to the main classifier \mathcal{M}^Y , we introduce an adversarial classifier \mathcal{M}^Z , which

takes X_S as the input and predicts the demographic attribute Z . During the adversarial training, the corrector attempts to help \mathcal{M}^Y make correct predictions while preventing \mathcal{M}^Z from predicting demographic attributes. To make this feasible, we use the gradient reversal technique (Ganin and Lempitsky, 2015), where we add a gradient-reversal layer between the weighted inputs X_S and the adversarial classifier \mathcal{M}^Z . The gradient-reversal layer has no effect on its downstream components (i.e., the adversarial classifier \mathcal{M}^Z). However, during back-propagation, the gradients that pass down through this layer to its upstream components (i.e., the corrector \mathcal{C}) are getting reversed. As a result, the corrector \mathcal{C} receives opposite gradients from \mathcal{M}^Z . The outputs of the \mathcal{M}^Y and \mathcal{M}^Z are used as signals to train the corrector such that it can upweight the words correlated with the main task label Y and downweight the words correlated with the demographic attribute Z . We set the adversarial classifier \mathcal{M}^Z with the same architecture as the main classifier \mathcal{M}^Y . The corrector \mathcal{C} has the same architecture as the explainer introduced in Section 3.

4.2 An Optimization Method for Debaised-TC

In this subsection, we discuss the optimization method for the proposed framework. We denote the parameters of \mathcal{M}^Y , \mathcal{M}^Z and \mathcal{C} as \mathbf{W}^Y , \mathbf{W}^Z and Θ , respectively. The optimization task is to jointly optimize the parameters of the classifiers, i.e., \mathbf{W}^Y and \mathbf{W}^Z , and the parameters of the corrector, i.e., Θ . We can view the optimization as an architecture search problem. Since our debiasing framework is end-to-end and differentiable, we develop an optimization method for our framework based on the differentiable architecture search (DARTS) technique (Liu et al., 2018). We update \mathcal{M}^Y , \mathcal{M}^Z by optimizing the training losses L_{train}^Y and L_{train}^Z on the training set and update Θ by optimizing the validation loss L_{val} on the validation set through gradient descent. We denote the cross-entropy losses for \mathcal{M}^Y and \mathcal{M}^Z as L^Y and L^Z , respectively. L_{train}^Y and L_{train}^Z indicate the cross-entropy losses L^Y and L^Z on the training set. L_{val} denotes the combined loss of the two cross-entropy losses $L = L^Y + L^Z$ on the validation set.

The goal of optimizing the corrector is to find optimal parameters Θ^* that minimizes the validation loss $L_{val}(\mathbf{W}^{Y*}, \mathbf{W}^{Z*}, \Theta)$, where the optimal parameters \mathbf{W}^{Y*} and \mathbf{W}^{Z*} are obtained by mini-

mizing the training losses as follows.

$$\begin{aligned}\mathbf{W}^{Y*} &= \arg \min_{\mathbf{W}^Y} L_{train}^Y(\mathbf{W}^Y, \Theta^*) \\ \mathbf{W}^{Z*} &= \arg \min_{\mathbf{W}^Z} L_{train}^Z(\mathbf{W}^Z, \Theta^*)\end{aligned}$$

The above goal forms a bi-level optimization problem (Maclaurin et al., 2015; Pham et al., 2018), where Θ is the upper-level variable and \mathbf{W}^Y and \mathbf{W}^Z are the lower-level variables:

$$\begin{aligned}\min_{\Theta} L_{val}(\mathbf{W}^{Y*}(\Theta), \mathbf{W}^{Z*}(\Theta), \Theta) \\ s.t. \mathbf{W}^{Y*}(\Theta) &= \arg \min_{\mathbf{W}^Y} L_{train}^Y(\mathbf{W}^Y, \Theta^*) \\ \mathbf{W}^{Z*}(\Theta) &= \arg \min_{\mathbf{W}^Z} L_{train}^Z(\mathbf{W}^Z, \Theta^*)\end{aligned}$$

Optimizing Θ is time-consuming due to the expensive inner optimization of \mathbf{W}^Y and \mathbf{W}^Z . Therefore, we leverage the approximation scheme as DARTS:

$$\begin{aligned}\nabla_{\Theta} L_{val}(\mathbf{W}^{Y*}(\Theta), \mathbf{W}^{Z*}(\Theta), \Theta) \\ \approx \nabla_{\Theta} L_{val}(\mathbf{W}^Y - \xi \nabla_{\mathbf{W}^Y} L_{train}^Y(\mathbf{W}^Y, \Theta), \\ \mathbf{W}^Z - \xi \nabla_{\mathbf{W}^Z} L_{train}^Z(\mathbf{W}^Z, \Theta), \Theta)\end{aligned}$$

where ξ is the learning rate for updating \mathbf{W}^Y and \mathbf{W}^Z . The approximation scheme estimates $\mathbf{W}^{Y*}(\Theta)$ and $\mathbf{W}^{Z*}(\Theta)$ by updating \mathbf{W}^Y and \mathbf{W}^Z for a single training step, which avoids the total optimization $\mathbf{W}^*(\Theta) = \arg \min_{\mathbf{W}} L_{train}(\mathbf{W}, \Theta^*)$ to the convergence. In our implementation, we apply first-order approximation with $\xi = 0$, which can even lead to more speed-up. Also, in our specific experiments, since the amount of validation data is limited, we build an augmented validation dataset $\mathcal{V}' = \mathcal{V} \cup \mathcal{T}$ combining the original validation set \mathcal{V} with the training set \mathcal{T} for optimizing Θ .

We present our DARTS-based optimization algorithm in Algorithm 1. In each iteration, we first update the corrector’s parameters based on the augmented validation set \mathcal{V}' (lines 2-3). Then, we collect a new mini-batch of training data (line 4). We generate the saliency scores $S = (s_1, s_2, \dots, s_n)$ for the training examples via the corrector with its current parameters (line 5). Next, we make predictions via the classifiers with their current parameters and X_S (line 6). Eventually, we update the parameters of the classifiers (line 7).

5 Experiment

In this section, we conduct experiments to evaluate our proposed debiasing framework. Through the

Algorithm 1: The DARTS-based optimization method for Debiased-TC.

- 1 **Input:** Training data $\mathcal{T} = \{X_i, Y_i, Z_i\}_{i=1}^{|\mathcal{T}|}$ and Validation data $\mathcal{V} = \{X_i, Y_i, Z_i\}_{i=1}^{|\mathcal{V}|}$
 - 2 **Output:** classifier parameters \mathbf{W}^{Y*} and \mathbf{W}^{Z*} ; and corrector parameters Θ^*
 - 3 Initialize \mathbf{W}^Y , \mathbf{W}^Z and Θ
 - 1: **while** not converged **do**
 - 2: Sample a mini-batch of validation data from $\mathcal{V}' = \mathcal{V} \cup \mathcal{T}$
 - 3: Update Θ by descending $\nabla_{\Theta} L_{val}(\mathbf{W}^Y - \xi \nabla_{\mathbf{W}^Y} L_{train}^Y(\mathbf{W}^Y, \Theta), \mathbf{W}^Z - \xi \nabla_{\mathbf{W}^Z} L_{train}^Z(\mathbf{W}^Z, \Theta), \Theta)$ ($\xi = 0$ for first-order approximation)
 - 4: Collect a mini-batch of training data from \mathcal{T}
 - 5: Generate S via the corrector with current parameters Θ
 - 6: Generate predictions via the classifiers with current parameters \mathbf{W}^Y , \mathbf{W}^Z and X_S
 - 7: Update \mathbf{W}^Y and \mathbf{W}^Z by descending $\nabla_{\mathbf{W}^Y} L_{train}^Y(\mathbf{W}^Y, \Theta)$ and $\nabla_{\mathbf{W}^Z} L_{train}^Z(\mathbf{W}^Z, \Theta)$
 - 8: **end while**
-

experiments, we try to answer two questions: 1) Does our framework effectively mitigate the implicit bias in various deep text classification models? and 2) Does our framework maintain the performance of the original models (without debiasing) while reducing the bias?

5.1 Baselines

In our experiments, we compare our proposed debiasing framework with two baselines. Since there is no established method for mitigating implicit bias, we adopt two debiasing methods designed for traditional explicit bias and adapt them for implicit bias.

Data Augmentation* (Data Aug) (Dixon et al., 2018). We manually balance the training data of two demographic groups by adding sufficient negative examples for Group I and positive examples for Group II. As a result, the ratio of positive and negative training examples for both groups is 1:1. As discussed in the introduction, obtaining additional labeled data from specific authors is very expensive. In this work, we seek to develop a bias mitigation methodology without extra data. *Since Data Aug introduces more training data, it's not fair to directly compare it with other debiasing methods that only utilize original training data (including our method). We include Data Aug as a special baseline for reference.*

Instance Weighting (Ins Weigh) (Zhang et al., 2020). We re-weight each training instance with a numerical weight $\frac{P(Y)}{P(Y|Z)}$ based on the label dis-

tribution for each demographic group to mitigate explicit bias. In this method, a random forest classifier is built to estimate the conditional distribution $P(Y|Z)$ and the marginal distribution $P(Y)$ is manually calculated.

5.2 Experimental Settings

We conduct our experiments for implicit bias mitigation on two representative base models: CNN (Kim, 2014) and RNN (Chung et al., 2014). We use the same datasets with manually designed proportions, as described in Section 2.2. The details of the base models, as well as the implementation details of the experiments, can be found in Appendix C.

5.3 Performance Comparison

We train the base models with our proposed debiasing framework as well as the baseline debiasing methods. We report the performance on the test set in terms of fairness and classification performance.

Fairness Evaluation. Table 3 shows the results for fairness evaluation metrics: false positive equality difference (FPED), false negative equality difference (FNED), and DPD. FPED/FNED indicates the absolute value of the difference between the false positive/negative rates of the two groups. We make the following observations. First, the base models attain high FPED, FNED, and DPD, which indicates the existence of significant implicit bias towards the authors of the texts. Ins Weigh seems ineffective in mitigating implicit bias since it only achieved comparable fairness scores with the base models. Note that not every example that belongs to a certain group necessarily results in bias towards that group. Thus, assigning a uniform weight for all examples with the same label Y and demographic attribute Z is not a proper way to reduce implicit bias. Third, both Data Aug and Debiased-TC can mitigate the implicit bias by achieving lower equality and demographic parity differences. However, compared to Data Aug, Debiased-TC has two advantages. First, Data Aug needs to add more training data while Debiased-TC does not. Debiased-TC can locate the main source of implicit bias by analyzing how it forms in a deep text classification model. Due to the proposed corrector model, it can make a classification model focus on the relevant features for predictions and discard the features that may lead to implicit bias. Second, Debiased-TC is more stable than Data Aug. For the sentiment classification task with race as the demographic

Table 3: Fairness Performance Comparison. Note that Data Aug is a special baseline for reference.

Task	Methods	CNN			RNN		
		FPED (%)	FNED (%)	DPD (%)	FPED (%)	FNED (%)	DPD (%)
Sentiment Race (DIAL)	Base Model	23.59	41.45	32.52	26.86	42.36	34.61
	Data Aug*	21.00*	3.88*	12.44*	19.84*	0.59*	10.22*
	Ins Weigh	25.47	41.43	33.45	26.86	42.36	34.61
	Debiased-TC	6.08	4.63	0.73	6.67	5.68	0.50
Mention Race (DIAL)	Base Model	32.73	17.58	25.16	30.44	17.55	24.00
	Data Aug*	1.31*	7.31*	3.00*	0.77*	7.91*	4.34*
	Ins Weigh	24.66	19.46	22.06	28.83	17.26	23.05
	Debiased-TC	3.61	2.40	0.61	4.97	1.07	1.95
Mention Gender (PAN16)	Base Model	11.60	9.95	10.78	10.62	8.33	9.47
	Data Aug*	0.84*	0.19*	0.32*	2.42*	0.72*	1.57*
	Ins Weigh	12.73	10.22	11.47	11.20	9.35	10.28
	Debiased-TC	3.95	3.04	3.49	5.41	3.73	4.57
Mention Age (PAN16)	Base Model	15.03	9.96	12.49	13.07	7.34	10.20
	Data Aug*	3.71*	1.59*	1.06*	0.17*	2.69*	1.26*
	Ins Weigh	16.53	8.71	12.62	13.24	7.94	10.59
	Debiased-TC	7.29	2.91	5.10	7.64	2.69	5.16
Hate Speech Race (MTC)	Base Model	78.56	37.22	57.89	81.51	28.50	55.01
	Data Aug*	88.81*	26.15*	57.48*	83.51*	22.73*	53.12*
	Ins Weigh	87.51	31.92	59.72	84.45	27.44	55.95
	Debiased-TC	75.97	17.08	46.53	74.56	18.85	46.70

attribute, the CNN and RNN classifiers trained on augmented data still result in high FPED and DPD scores. This suggests that balancing the training data cannot always mitigate implicit bias. In fact, only training examples with demographic language features can contribute to the implicit bias. Since some texts in the training set do not contain any language features belonging to a demographic group, they do not help balance the data.

Text Classification Performance Evaluation.

The prediction performance of the text classification models trained under various debiasing methods is shown in Table 4, where we report the accuracy and F1 scores. First, it is not surprising to see that Data Aug achieves the best performances, since the data augmentation technique introduces more training data. It’s not fair to directly compare it with other debiasing methods that only utilize original training data. Second, in most cases, our method achieves comparable or even better performance than the original base models. As we verified before, the implicit bias of a text classification model is caused by the fact that it learns a wrong correlation between labels and demographic language features. Debiased-TC corrects the model’s selection of language features for predictions and thereby improves its performance on the classification task.

In conclusion, our proposed debiasing framework significantly mitigates the implicit bias, while maintaining or even slightly improving the classification performance.

cation performance.

6 Related Work

Fairness in Machine Learning. With the wide spread of the machine learning (ML) applications in our daily lives, bias and fairness issues in them are drawing increasing attention from the community. Researches are conducted to detect and mitigate the bias in ML models on various tasks. Specifically, studies investigate how algorithms can be biased in classification (Kamiran and Calders, 2009; Chouldechova, 2017), regression (Berk et al., 2017; Agarwal et al., 2019), and clustering tasks (Backurs et al., 2019; Chen et al., 2019). In the domain of computer vision, researchers show that ML-based face recognition (Buolamwini and Gebru, 2018) and object detection (Ryu et al., 2017) models perform unfairly for different demographic groups. Besides, a lot of works examine the bias in language related tasks, including word embedding (Bolukbasi et al., 2016), coreference resolution (Zhao et al., 2018), machine translation (Prates et al., 2019) and dialogue generation (Liu et al., 2020a,b), etc. Moreover, some recent studies also explore the relationship between the fairness of an ML model and its other properties, such as robustness (Xu et al., 2020; Nanda et al., 2021) and privacy (Cummings et al., 2019).

Fairness in Text Classification. In this work, we focus on the fairness issues in the text classification task. In this task, Dixon et al. (2018)

Table 4: Text Classification Performance Comparison (%). Note that Data Aug is a special baseline for reference.

Methods	Sentiment/Race (DIAL)		Mention/Race (DIAL)		Mention/Gender (PAN16)		Mention/Age (PAN16)		Hate Speech/Race (MTC)	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
CNN										
Base Model	61.40	60.03	70.77	71.65	81.93	81.94	80.57	80.17	64.10	65.86
Data Aug*	67.58*	71.53*	76.42*	76.03*	84.11*	84.31*	84.08*	84.36*	66.96*	71.10*
Ins Weigh	61.06	60.36	71.62	69.66	81.86	81.85	80.70	81.05	65.25	68.73
Debiased-TC	63.60	66.58	73.15	71.84	81.67	82.01	80.41	79.68	69.14	72.69
RNN										
Base Model	61.23	61.53	72.97	73.68	83.46	83.40	82.78	82.43	66.31	69.57
Data Aug*	67.82*	69.35*	78.42*	77.26*	86.25*	86.05*	86.12*	85.68*	68.55*	72.37*
Ins Weigh	61.23	61.53	73.37	73.79	83.46	83.32	82.80	82.58	67.26	70.94
Debiased-TC	63.68	66.70	74.05	73.41	81.81	81.51	80.21	79.17	66.76	70.76

demonstrate that the source of unintended bias in models is the imbalance of training data, and they provide a debiasing method, which introduces new data to balance the training data. In (Park et al., 2018), gender bias is measured on abusive language detection models, and the effects of different pre-trained word embeddings and model architectures are analyzed. By considering the various ways that a classifier’s score distribution can vary across designated groups, a suite of threshold-agnostic metrics is introduced in (Borkan et al., 2019), which provides a nuanced view of unintended bias. Furthermore, the work (Zhang et al., 2020) proposes to debias text classification models using instance weighting, i.e., different weights are assigned to the training samples involving different demographic groups. The works discussed above focus on explicit bias, where the demographic attributes are explicitly expressed in the text. However, works studying implicit bias are rather limited. Huang et al. (2020) introduce the first multilingual hate speech dataset with inferred author demographic attributes. Through experiments on this dataset, they show that popular text classifiers can learn the bias towards the demographic attribute of the author. But this work doesn’t discuss how the bias is produced, and no debiasing method is provided.

7 Conclusion

In this paper, we demonstrate that a text classifier with implicit bias makes predictions based on language features correlated with demographic groups of authors, and propose a novel learning framework Debiased-TC to mitigate such implicit bias. The experimental results show that Debiased-TC significantly mitigates implicit bias, and maintains or even improves the text classification performance of the original models. In the future, we will inves-

tigate implicit bias in other NLP applications.

Acknowledgments

Haochen Liu, Wei Jin, Hamid Karimi and Jiliang Tang are supported by the National Science Foundation (NSF) under grant number CNS1815636, IIS1928278, IIS1714741, IIS1845081, IIS1907704, and IIS1955285. Zitao Liu is supported by Beijing Nova Program (Z201100006820068) from Beijing Municipal Science & Technology Commission.

References

- Alekh Agarwal, Miroslav Dudik, and Zhiwei Steven Wu. 2019. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, pages 120–129. PMLR.
- Arturs Backurs, Piotr Indyk, Krzysztof Onak, Baruch Schieber, Ali Vakilian, and Tal Wagner. 2019. Scalable fair clustering. In *International Conference on Machine Learning*, pages 405–413. PMLR.
- Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2017. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476.
- Su Lin Blodgett, Lisa Green, and Brendan O’Connor. 2016. Demographic dialectal variation in social media: A case study of african-american english. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130.

- 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.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In Conference on fairness, accountability and transparency, pages 77–91. PMLR.
- Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. 2018. Learning to explain: An information-theoretic perspective on model interpretation. In International Conference on Machine Learning, pages 883–892. PMLR.
- Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally fair clustering. In International Conference on Machine Learning, pages 1032–1041. PMLR.
- Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. Big data, 5(2):153–163.
- Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In NIPS 2014 Workshop on Deep Learning, December 2014.
- Florian Coulmas. 2013. Sociolinguistics: The study of speakers’ choices. Cambridge University Press.
- Thomas M Cover. 1999. Elements of information theory. John Wiley & Sons.
- Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. 2019. On the compatibility of privacy and fairness. In Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization, pages 309–315.
- 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.
- 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.
- Yanai Elazar and Yoav Goldberg. 2018. Adversarial removal of demographic attributes from text data. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 11–21.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In International conference on machine learning, pages 1180–1189. PMLR.
- Xiaolei Huang, Linzi Xing, Franck Dernoncourt, and Michael Paul. 2020. Multilingual twitter corpus and baselines for evaluating demographic bias in hate speech recognition. In Proceedings of The 12th Language Resources and Evaluation Conference, pages 1440–1448.
- Faisal Kamiran and Toon Calders. 2009. Classifying without discriminating. In 2009 2nd International Conference on Computer, Control and Communication, pages 1–6. IEEE.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1746–1751. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In ICLR (Poster).
- Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. In Advances in neural information processing systems, pages 4066–4076.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. Darts: Differentiable architecture search. In International Conference on Learning Representations.
- Haochen Liu, Jamell Dacon, Wenqi Fan, Hui Liu, Zitao Liu, and Jiliang Tang. 2020a. Does gender matter? towards fairness in dialogue systems. In Proceedings of the 28th International Conference on Computational Linguistics, pages 4403–4416.
- Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zitao Liu, and Jiliang Tang. 2020b. Mitigating gender bias for neural dialogue generation with adversarial learning. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 893–903.
- Dougal Maclaurin, David Duvenaud, and Ryan Adams. 2015. Gradient-based hyperparameter optimization through reversible learning. In International Conference on Machine Learning, pages 2113–2122.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. End-to-end bias mitigation by modelling biases in corpora. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 8706–8716.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2019. A survey on bias and fairness in machine learning. arXiv preprint arXiv:1908.09635.

Vedant Nanda, Samuel Dooley, Sahil Singla, Soheil Feizi, and John P Dickerson. 2021. Fairness through robustness: Investigating robustness disparity in deep learning. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pages 466–477.

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.

Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. 2018. Efficient neural architecture search via parameter sharing. In ICML.

Marcelo OR Prates, Pedro H Avelar, and Luis C Lamb. 2019. Assessing gender bias in machine translation: a case study with google translate. Neural Computing and Applications, pages 1–19.

Daniel PreoŃiuc-Pietro and Lyle Ungar. 2018. User-level race and ethnicity predictors from twitter text. In Proceedings of the 27th International Conference on Computational Linguistics, pages 1534–1545.

Francisco Rangel, Paolo Rosso, Ben Verhoeven, Walter Daelemans, Martin Potthast, and Benno Stein. 2016. Overview of the 4th author profiling task at pan 2016: cross-genre evaluations. Working Notes Papers of the CLEF, 2016:750–784.

Hee Jung Ryu, Margaret Mitchell, and Hartwig Adam. 2017. Improving smiling detection with race and gender diversity. arXiv preprint arXiv:1712.00193, 1(2):7.

Deven Santosh Shah, H Andrew Schwartz, and Dirk Hovy. 2020. Predictive biases in natural language processing models: A conceptual framework and overview. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 5248–5264.

Han Xu, Xiaorui Liu, Yaxin Li, and Jiliang Tang. 2020. To be robust or to be fair: Towards fairness in adversarial training. arXiv preprint arXiv:2010.06121.

Guanhua Zhang, Bing Bai, Junqi Zhang, Kun Bai, Conghui Zhu, and Tiejun Zhao. 2020. Demographics should not be the reason of toxicity: Mitigating discrimination in text classifications with instance weighting. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 4134–4145.

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 (Short Papers), pages 15–20.

A Details of the Datasets

In this section, we describe the statistical information and the annotation process of the three datasets DIAL, PAN16, and MTC used in our experiments. The datasets DIAL and PAN16 can be downloaded from the link ². The dataset MTC can be downloaded from the link ³.

A.1 Data Statistics

The statistics of the datasets DIAL, PAN16, and MTC are shown in Table 5. In the table, the “task” section shows the text classification tasks included in a dataset. “Sentiment” is short for sentiment analysis. “Mention” is short for mention detection. “Hate Speech” is short for hate speech detection. “Demog.” indicates the demographic attribute of the tweet authors collected in a dataset. The “Size” section shows the total number of instances in a dataset. Each instance is a tweet text. The “Avg.Len.” section shows the average number of words in one instance in a dataset.

Table 5: Statistics of the datasets.

Dataset	Task	Demog.	Size	Avg.Len.
DIAL	Sentiment	Race	317,151	11.20
	Mention	Race	400,000	10.56
PAN16	Mention	Gender	175,871	14.64
	Mention	Age	175,471	14.55
MTC	Hate Speech	Race	47,627	19.60

A.2 Data Annotation

The DIAL dataset is annotated based on the dialectal tweet corpus (Blodgett et al., 2016), which contains 59.2 million tweets from 2.8 million users. The race attribute is annotated by an automated probabilistic inference method based on the geolocation information of the user and the tweet text. Given that geolocation information (residence) is highly associated with the race of a user, the model can make accurate predictions. To further ensure the accuracy, DIAL only keeps annotations with confidence above 80%.

²<https://github.com/yanaiela/demog-text-removal>

³https://github.com/xiaoleihuang/Multilingual_Fairness_LREC

The PAN16 dataset (Rangel et al., 2016) contain 436 Twitter users, each of which has up to 1,000 tweets. The age and gender of the users are manually annotated by referring to their LinkedIn profiles. Specifically, annotators judge the gender based on the user’s name and profile photo. The age is inferred based on the user’s birth date or degree starting date.

The MTC dataset (Huang et al., 2020) is annotated based on 7 published Twitter hate speech datasets in five languages. The dataset contains user demographic information such as race, gender, age, and country. We only focus on the English corpus and the attribute race in our experiments. The race of a user is inferred by the computer vision API, Face++⁴, based on the profile photo.

B Optimization of the Explainer

We train the explainer \mathcal{E} by maximizing the mutual information between the response variable Y and the selected features X_S . The optimization problem can be formulated as:

$$\begin{aligned} \max_{\mathcal{E}} \quad & I(X_S; Y) \\ \text{s.t.} \quad & S \sim P_{\mathcal{E}}(S|X) \end{aligned} \quad (1)$$

where

$$\begin{aligned} I(X_S, Y) &= \mathbb{E} \left[\log \frac{P(X_S, Y)}{P(X_S)P(Y)} \right] \\ &= \mathbb{E} \left[\log \frac{P_{\mathcal{M}}(Y|X_S)}{P(Y)} \right] \\ &\propto \mathbb{E} \left[\log P_{\mathcal{M}}(Y|X_S) \right] \\ &= \mathbb{E}_X \mathbb{E}_{S|X} \mathbb{E}_{Y|X_S} \left[\log P_{\mathcal{M}}(Y|X_S) \right] \end{aligned}$$

Solving the optimization problem in Eq. (1) is equivalent to finding an explainer \mathcal{E} satisfying the following:

$$\max_{\mathcal{E}} P_{\mathcal{M}}(Y|X_S) \quad \text{s.t.} \quad S \sim P_{\mathcal{E}}(S|X).$$

Hence, we train the explainer \mathcal{E} by optimizing $P_{\mathcal{M}}(Y|X_S)$ with the parameters of the classification model \mathcal{M} fixed. In our implementation, we adopt the cross-entropy loss for training, as we do when we train the classification model \mathcal{M} .

⁴<https://www.faceplusplus.com/>

C Implementation Details

C.1 Details of Base Models

In the base model CNN, we use 100 filters with three different kernel sizes (3, 4, and 5) in the convolution layer, where we use a Rectified Linear Unit (ReLU) as the non-linear activation function. Each obtained feature map is processed by a max-pooling layer. Then, the features are concatenated and fed into a linear prediction layer to get the final predictions. A dropout with a rate of 0.3 is applied before the linear prediction layer.

For the base model RNN, we use a one-layer unidirectional RNN with Gated Recurrent Units (GRU). The hidden size is set to 300. The last hidden state of the RNN is fed into a linear prediction layer to get the final predictions. We apply a dropout with a rate of 0.2 before the linear prediction layer.

C.2 Details of Experimental Settings

For the text classifiers, we use randomly initialized word embeddings with a size of 300. All the models are trained by an Adam optimizer (Kingma and Ba, 2015) with an initial learning rate of 0.001. We apply gradient clipping with a clip-value of 0.25 to prevent the exploding gradient problem. The batch size is set to 64. For the base model and the baseline methods, when the prediction accuracy of the validation data doesn’t improve for 5 consecutive epochs, the training is terminated, and we pick the model with the best performance on the validation set. Our model utilizes the validation data for training. To avoid it overfitting the validation data, we don’t select the model based on its performance on the validation set. Instead, we train the model for a fixed number of epochs (5 epochs, the same for all the three datasets) and evaluate the obtained model.