

©Copyright 2020

Xuhui Zhou

# Challenges in Automated Debiasing for Toxic Language Detection

Xuhui Zhou

A thesis  
submitted in partial fulfillment of the  
requirements for the degree of

Master of Science

University of Washington

2020

Committee:

Noah A. Smith

Shane Steinert-Threlkeld

Program Authorized to Offer Degree:  
Linguistics

University of Washington

**Abstract**

Challenges in Automated Debiasing for Toxic Language Detection

Xuhui Zhou

Chair of the Supervisory Committee:

Noah A. Smith

Paul G. Allen School of Computer Science & Engineering

**Warning:** *this thesis contains content that may be offensive or upsetting.*

Biased associations have been a challenge in the development of classifiers for detecting toxic language, hindering both fairness and accuracy. As potential solutions, we investigate recently introduced debiasing methods for text classification datasets and models, as applied to toxic language detection. Our focus is on lexical (e.g., swear words, slurs, identity mentions) and dialectal markers (specifically African American English). Our comprehensive experiments establish that existing methods are limited in their ability to prevent biased behavior in current toxicity detectors. We then propose an automatic, dialect-aware data correction method, as a proof-of-concept study. Despite the use of synthetic labels, this method reduces dialectal associations with toxicity. Overall, our findings show that debiasing a model trained on biased toxic language data is not as effective as simply relabeling the data to remove existing biases. This work is published at EACL 2021: The 16th Conference of the European Chapter of the Association for Computational Linguistics.

# TABLE OF CONTENTS

	Page
List of Figures . . . . .	iii
List of Tables . . . . .	iv
Chapter 1: Introduction . . . . .	1
1.1 Contributions . . . . .	2
Chapter 2: Related Work . . . . .	5
2.1 Toxicity Datasets Creation . . . . .	5
2.2 Toxicity Classification . . . . .	6
2.3 Bias in Toxicity Datasets / Models . . . . .	6
2.4 Debiasing Toxicity Detection . . . . .	7
2.5 Other General Debiasing Methods . . . . .	7
Chapter 3: Biases in Toxic Language Detection . . . . .	8
3.1 Lexical Biases (TOXTRIG) . . . . .	9
3.2 Dialectal Biases (AAE) . . . . .	10
3.3 Dataset for Toxic Language Detection . . . . .	10
Chapter 4: Debiasing Methods . . . . .	12
4.1 Debaised Training for Pre-Defined Toxicity Biases . . . . .	12
4.2 Data Filtering for Spurious Biases . . . . .	13
Chapter 5: Experiments: Lexical Biases . . . . .	15
5.1 Evaluation Framework . . . . .	15
5.2 Results for Lexical Bias Reduction . . . . .	16
5.3 Qualitative Analysis . . . . .	18
5.4 Adversarial Evaluation: ONI-Adv . . . . .	19

Chapter 6:	Experiments: Dialectal and Racial Biases . . . . .	21
6.1	Dialectal Biases . . . . .	21
6.2	Racial Biases . . . . .	21
Chapter 7:	Towards Data Relabeling . . . . .	25
Chapter 8:	Conclusion . . . . .	28
Appendix A:	Further Details for Models . . . . .	41
A.1	Model Debiasing . . . . .	41
A.2	Data Filtering . . . . .	41
A.3	Training Settings . . . . .	41
A.4	Prediction Combining with Bias-only Model . . . . .	42
Appendix B:	Alternative Dataset of Toxic Language . . . . .	43
Appendix C:	Few-shot AAE-to-WAE Translation . . . . .	46

## LIST OF FIGURES

Figure Number	Page
1.1 Lexical items and dialect markers cause problematic behavior for toxic language detection systems such as the widely used PerspectiveAPI. In the top two example pairs, statements with minority identity mentions and swear words used inoffensively are flagged as toxic, but majority identity mentions or offensive statements without overt swearing are missed. The bottom pair shows dialect-based racial bias for two inoffensive greetings, where markers of African American English (AAE) trigger the toxicity detector. . . . .	4
5.1 Challenge set evaluation for lexical biases, comparing all debiasing methods with baselines, using the ONI-Adv test set. <b>Takeaway:</b> $F_1(\uparrow)$ measures show that all models perform poorly at identifying toxic text not containing overtly lexical cues of toxicity. In general, debiased training approaches outperform the original model on this challenge set, while data filtering is not as effective. . . . .	20

## LIST OF TABLES

Table Number		Page
5.1	Lexical associations between toxicity and TOXTRIG mentions in the original dataset (Founta et al., 2018) and various filtered counterparts. Random, AFLite, and DataMaps all contain only 33% of the original data after filtering. Lower Pearson $R$ correlation value indicates less superficial patterns in the dataset, i.e., less bias. <b>Takeaway:</b> The hard and ambiguous subsets given by DataMaps contain the lowest amount of lexical associations, indicated in boldface. . . . .	16
5.2	Evaluation of lexical bias removal for all debiasing methods on the Founta et al. (2018) test set. Results show the mean and s.d. (subscript) of accuracy and $F_1$ across 3 runs, as well as $F_1$ and false positive rate exclusively on test examples containing specific TOXTRIG mentions—NOI, OI and ONI, along with the number of examples in each category. The lower the FPR*, the less the model infers lexical associations for toxicity. The first block shows debiased training approaches, along with the vanilla classifier, which are trained on the full dataset. The second block shows data filtering approaches, all trained on only 33% of the training data. Best performance in each block is boldfaced. <b>Takeaway:</b> While data filtering approaches achieve overall higher performance, debiased training approaches perform better on lexical bias reduction, in aggregate. . . . .	17
5.3	Examples of test set tweets with their gold-standard annotations and predictions from models trained on DataMaps-Hard (DM-Hard) and DataMaps-Easy (DM-Easy) subsets. 🦠 denotes tweets with toxic labels, and 🍃 represents non-toxic labels. We anonymize the usernames to protect user privacy. . . . .	18

6.1	Dialectal bias evaluation for all debiasing methods (§6), as well as the relabeling approach (§7) on the Founta et al. (2018) test set. We report $F_1$ and the false positive rate with respect to tweets in AAE ( $FPR_{AAE}$ ), reflecting dialectal bias (lower is less biased), showing mean and s.d. (subscript) across 3 runs. (Top Block) Debaised training approaches, along with the vanilla classifier, are all trained on the full dataset. (Middle Block) Random, AFLite and DataMaps all are trained on only 33% of the training data. Best performance for each training set size is in boldface. <b>Takeaway:</b> Both debiasing approaches improve performance over baselines, with DataMaps-Hard proving the most effective at debiasing. (Bottom Block) AAE-relabeling results in a model which despite following a noisy process yields even larger improvements for dialectal debiasing. . . . .	23
6.2	Racial disparity in toxicity prediction reported on Preoțiu-Pietro and Ungar (2018). <b>W-Tox.</b> indicates % of white users’ tweets being flagged as toxic, <b>AA-Tox.</b> indicates % of African American users’ tweets being flagged as toxic, $\Delta$ refers to the difference between AA-Tox. and W-Tox., and <b>AA/W</b> refers to the ratio between AA-Tox. and W-Tox. <b>Takeaway:</b> Methods generally fail in debiasing on this OOD test set except the relabeling approach shows some benefit. . . . .	24
7.1	Examples of AAE tweets with their GPT-3 based WAE translation, and original gold standard and new annotations based on AAE-relabeled. For the first three tweets, the (biased) gold labels are changed by models predicting the new labels on their WAE translations. 🦠 indicates presence of toxicity, and 🍃 represents non-toxic. We anonymize the usernames to protect user privacy. . . . .	26
B.1	Lexical and dialectal associations between toxicity in the original dataset (Davidson et al., 2017) and various filtered counterparts. Random, AFLite, and DataMaps all contain only 50% of the original data after filtering. (We could not perform downsampling on these datasets due to their heavily skewed label distribution.) Lower Pearson $R$ correlation value indicates less superficial patterns in the dataset, thus are less biased. The easy subset gives the best results here are due to its severe imbalanced label distribution. . . . .	44
B.2	Lexical bias removal evaluation for debiasing methods. Original refers to the model trained over the full training set. The test set is further categorized into tweets that contained relevant TOXTRIG words. $F_1$ indicates models’ performance while the false positive rate ( $FPR_*$ ) reflects models’ bias. The lower the $FPR_*$ is, the less biased the model tend to be. . . . .	44

B.3 Dialectal bias evaluation for all debiasing methods, on both in-distribution test set as well as out-of-distribution dialect and race priming test sets. In addition to accuracy and  $F_1$ , we report the false positive rate with respect to tweets in AAE ( $FPR_{AAE}$ ), reflecting dialectal bias (lower is less debiased). Each method is based on a RoBERTa-large classifier. . . . . 45

## **ACKNOWLEDGMENTS**

I would like to thank my advisor Noah Smith for supervising me throughout my graduate studies and Maarten Sap, Swabha Swayamdipta as well as Yejin Choi for advising me in this project. I would also like to thank Shane Steinert-Threlkeld for being my thesis reader.

## Chapter 1

### INTRODUCTION

Current hate speech or toxic language detection<sup>1</sup> systems exhibit problematic and discriminatory behavior that causes them to have disparate negative impact on minority populations (Yasin, 2018; Guynn, 2020; Kim et al., 2020; Dias Oliva et al., 2020). Tweets simply containing a minority identity mention are commonly flagged as toxic by current systems, in contrast to those containing majority identity mentions, as illustrated in Figure 1.1.

At the core of the issue are *dataset biases*, i.e., spurious correlations between surface patterns and annotated toxicity labels (§3), which stem from the data creation process (Sap et al., 2019). Previous work has outlined two such biases for hate speech datasets (both shown in Figure 1.1): *lexical bias* which associates toxicity with the presence of certain words (e.g., profanities, identity mentions) (Dixon et al., 2018; Dinan et al., 2019) and *dialectal bias*, where toxicity is correlated with surface markers of African American English (AAE) (Davidson et al., 2019; Sap et al., 2019). When trained on biased datasets, models acquire *and* exacerbate these biases (e.g., flagging text by Black authors as more toxic than by white authors) (Sap et al., 2019; Zhang et al., 2018).

Concurrently, there has been elevated interest in developing *debiasing methods* for standard natural language understanding (NLU) tasks, i.e., methods that aim to decrease over-reliance on spurious correlations in NLU models (Clark et al., 2019; He et al., 2019; Karimi Mahabadi et al., 2020; Bras et al., 2020). This raises a natural question: *are current debiasing approaches effective for mitigating biases specific to toxic language detection?*

In this work, we address the above question by investigating two classes of debiasing approaches to mitigate lexical and dialectal biases: one that employs additional training objectives for bias removal, and another that filters training instances likely exhibiting spurious biases (§4).

---

<sup>1</sup>We use *hate speech* and *toxic language* interchangeably in this work, though their definitions do not perfectly align.

Through comprehensive experiments, we show that both approaches face major challenges in mitigating biases from a model trained on a biased dataset (in our case, the dataset from Founta et al. (2018)) for toxic language detection. While data filtering results in reduced bias associations in the data, *models* trained on filtered datasets still pick up on lexical (§5) and dialectal biases (§6). We find that dialectal biases are particularly challenging to address, as has also been shown by Xia et al. (2020). “Debiased” models still disproportionately flag text in certain dialects as toxic. Notably, mitigating dialectal bias through current debiasing methods does not mitigate a model’s propensity to label tweets by Black authors as more toxic than by white authors.

We additionally explore an alternative proof-of-concept study—relabeling supposedly toxic training instances whose automatic translations into a majority dialect are deemed non-toxic by the classifier. To this end, we create a synthetic dataset via few-shot dialect translation system built with GPT-3 (Brown et al., 2020). While only an illustrative solution, it nevertheless takes into account the dialectal context of the tweet, resulting in a model less prone to dialectal and racial biases (§7). Overall, our findings indicate that debiasing a model already trained on biased toxic language data can be challenging, compared to relabeling the data to remove existing biases. Our code and data are publicly available on Github.<sup>2</sup>

## 1.1 Contributions

This thesis contributes the following:

- We apply state-of-the-art debiasing methods to toxic detection, where those methods’ debiasing effects haven’t been tested before.
- Our systematic experiments and analysis suggest that current debiasing methods, although demonstrating a certain degree of debiasing, are limited in dealing with such subtle and noisy task.

---

<sup>2</sup>[https://github.com/XuhuiZhou/Toxic\\_Debias](https://github.com/XuhuiZhou/Toxic_Debias)

- We show an illustrative automatic data relabeling approach, which shows the most promising debiasing effects. This provides more insights for future efforts in debiasing toxicity detection.

The remainder of this thesis is organized as follows:

- **Chapter 2** introduces related work for this thesis.
- **Chapter 3** introduces the biases currently discovered in toxic language detection systems.
- **Chapter 4** introduces the debiasing methods using in this thesis.
- **Chapter 5** reports debiasing results for lexical biases.
- **Chapter 6** reports debiasing results for dialectal biases.
- **Chapter 7** explores automatically relabeling technique for debiasing toxic language system.
- **Chapter 8** draws a conclusion.

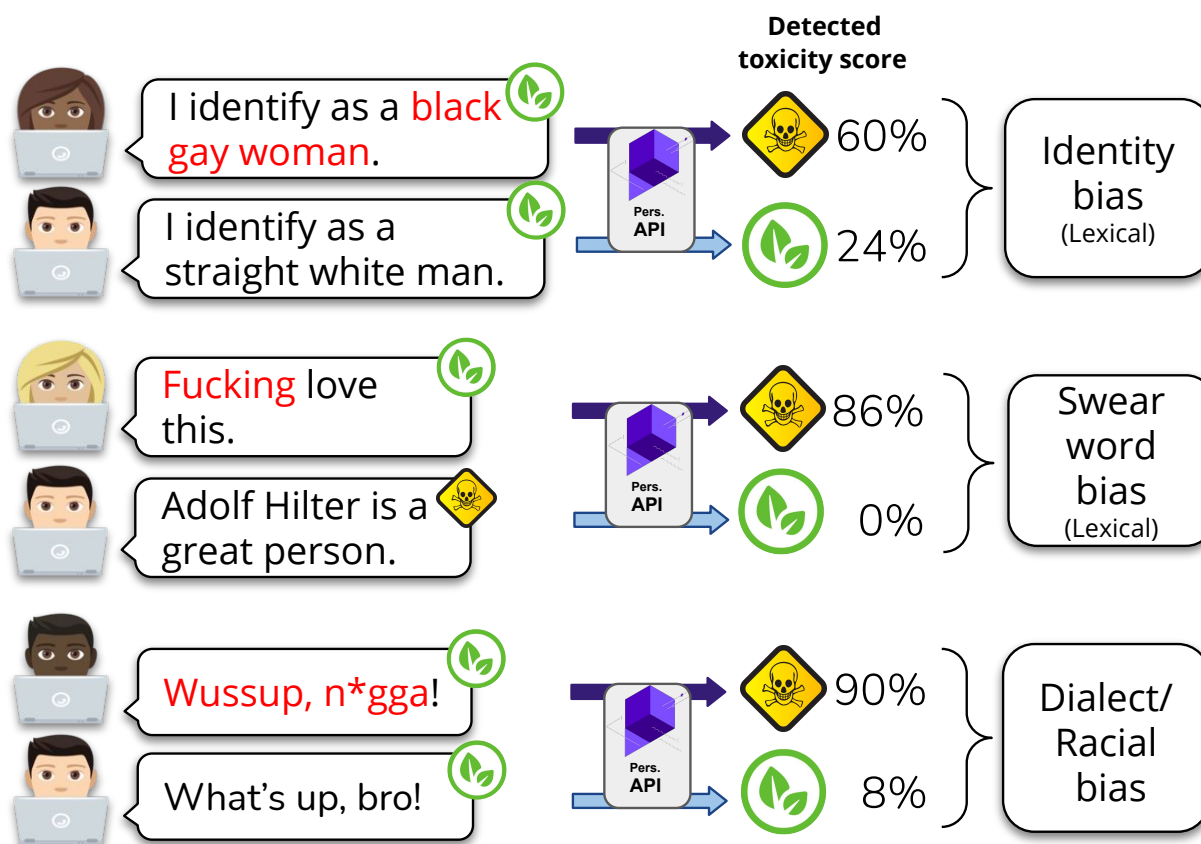


Figure 1.1: Lexical items and dialect markers cause problematic behavior for toxic language detection systems such as the widely used PerspectiveAPI. In the top two example pairs, statements with minority identity mentions and swear words used inoffensively are flagged as toxic, but majority identity mentions or offensive statements without overt swearing are missed. The bottom pair shows dialect-based racial bias for two inoffensive greetings, where markers of African American English (AAE) trigger the toxicity detector.

## Chapter 2

### RELATED WORK

Toxic language includes hate speech and offensive language, which is a big issue on the internet. A lot of efforts from companies and NLP people are happening. Many large social media platforms including Facebook, YouTube, and Instagram have deployed machine learning models on toxic language detection.<sup>1</sup> Kaggle hosted two competitions about toxic language detection.<sup>2,3</sup> Researchers continuously collect datasets (Davidson et al., 2017; Founta et al., 2018; Sap et al., 2019) and introduce models (Wulczyn et al., 2017; Noever, 2018) to tackle the toxic language in social media. Current systems have claimed satisfactory performance in detecting toxic language, typically, the first place algorithm Toxic Crusaders on Kaggle leader board of Toxic Comment Classification Challenge (Wulczyn et al., 2017) reached 98% accuracy score.

#### 2.1 Toxicity Datasets Creation

Many datasets were introduced to facilitate the automatic detection of toxic language. Each of them introduces different concepts in order to get a more fine-grained picture of toxic language. Davidson et al. (2017) first distinguish the difference between offensive language and hate speech, which they describe as *language that is used to express hatred towards a targeted group or is intended to be derogatory, to humiliate, or to insult the members of the group*. Founta et al. (2018) further classify toxic language as offensive, abusive, hateful, and spam, four sub-categories. To facilitate the study of detailed reasoning of toxic language, Sap et al. (2020) introduce Social Bias Frames.

---

<sup>1</sup><https://ai.facebook.com/blog/community-standards-report/>

<sup>2</sup><https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge>

<sup>3</sup><https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>

Those datasets may come from different sources including Twitter (Founta et al., 2018; Chatzakou et al., 2017; Golbeck et al., 2017; Waseem and Hovy, 2016), Facebook comments (Kumar et al., 2018), and Wiki comments (Wulczyn et al., 2017). Note that the sources all belong to the social media domain, which is reasonable to collect data from multiple sources to get a more comprehensive dataset. For example, Sap et al. (2020) annotate Social Bias Frames from Twitter and Reddit.

## **2.2 Toxicity Classification**

With the recent improvement of NLP techniques, many models have been proposed to tackle the problem. Those models include traditional machine learning approach such as logistic regression (Waseem and Hovy, 2016; Wulczyn et al., 2017; Davidson et al., 2017), and SVMs (Noever, 2018). However, the best performance usually comes from deep learning models (Badjatiya et al., 2017), such as convolutional neural network (CNN) (Pavlopoulos et al., 2017) and long short term memory network (LSTM) (Kumar et al., 2018).

While previous work in toxic content classification has focused on constructing classifiers that can flag toxic content with a high degree of accuracy, Kurita et al. (2019) utilize BERT (Devlin et al., 2019) to create models that are robust to character-level adversarial attack. Sap et al. (2020) apply GPT-2 (Radford et al., 2019) to further generate the reasons behind classification decisions.

## **2.3 Bias in Toxicity Datasets / Models**

Researchers have become aware that toxic language datasets and toxic language detection models have biases. Specifically, the bias here refers to lexical-based bias, including dialectal bias (Sap et al., 2019; Davidson et al., 2019) and identity-mentioned bias (Park et al., 2018).

Sap et al. (2019) investigate the racial bias in annotating toxic language datasets, and further show the models trained on these datasets amplify the bias against minority groups. Davidson et al. (2019) examine racial bias among five annotated Twitter dataset and show evidence of systematic racial bias in all datasets.

Park et al. (2018) show that abusive language detection models tend to be biased towards different gender identity. For example, *You are a good woman* is predicted as sexist by a model trained on certain dataset. Other identity-based bias also exists in current toxic language detection models; for example, *I am a gay man* is considered as highly toxic by some of the systems even if it is neutral.

## **2.4 Debiasing Toxicity Detection**

As the popularity of hate speech and toxic language detection systems has grown, several biases have been found in dataset and models, spurring various debiasing efforts to mitigate these individual biases (e.g., gender bias, racial bias) (Park et al., 2018; Sap et al., 2019; Davidson et al., 2019). Some work tackles identity-based biases, e.g., using data re-balancing (Dixon et al., 2018), or adversarial feature learning (Vaidya et al., 2019). Less work has tackled racial or dialectal bias. Notably, Xia et al. (2020) use adversarial training to prevent the model from associating toxicity with AAE, showing only small improvements in fairness. Based on those results, we do not explore adversarial methods, opting instead for ensemble-based methods of predefined bias reduction. In contemporary work, Mozafari et al. (2020) use a re-weighting mechanism, which shows some effects in debiasing racial bias. We leave it for future work to evaluate this method in our setting. In contrast to all previous work, our experiments also measure the effectiveness of bias-agnostic methods.

## **2.5 Other General Debiasing Methods**

Several approaches for debiasing NLU tasks have been proposed lately. Some approaches rely on adversarial training to remove protected attributes (e.g. gender or race), from a model’s internal representations (Zhang et al., 2018; Wang et al., 2019; Xia et al., 2020). Other approaches include confidence regularization (Utama et al., 2020), as well as other product of expert approaches (He et al., 2019; Karimi Mahabadi et al., 2020) similar to the debiased training approach from Clark et al. (2019), which is the only debiased training we employ due to its relatively strong performance.

## Chapter 3

### **BIASES IN TOXIC LANGUAGE DETECTION**

We test the use of debiasing<sup>1</sup> methods for the task of toxic language detection, which aims to flag rude, offensive, hateful, or toxic language on the internet, with the goal of moderating online communities (Roberts, 2019; Vidgen et al., 2019). This task differs in several ways from the natural language understanding (NLU) tasks that debiasing methods have been successful on, such as textual entailment (e.g., SNLI, MNLI) (Bowman et al., 2015; Williams et al., 2018) or reading comprehension (e.g., SQuAD) (Rajpurkar et al., 2016). First, compared to these NLU tasks where there is one correct label, the toxicity of language is inherently more nuanced, subjective, and contextual, which causes toxic language datasets to have lower agreement in general (Ross et al., 2017). Second, the dataset biases in NLU are predominantly artifacts introduced during data creation (e.g., negations, exaggerations) (Schwartz et al., 2017; Gururangan et al., 2018), whereas those in toxic language detection are grounded in the social dynamics of the world (Spears, 1998; Technau, 2018). For example, viewing AAE as a more toxic or less proper variety of English is a form of linguistic discrimination that upholds racial hierarchies in the United States (Rosa and Flores, 2017).

In this work, we consider two broad categories of toxic language dataset biases: lexical (§3.1) and dialectal (§3.2). Our experiments focus on a single, widely used dataset (§3.3) from Founta et al. (2018).

---

<sup>1</sup>Our definition of “bias” is specific to the social biases in toxic language detection datasets, grounded as lexical and dialectal biases; see Blodgett et al. (2020) for a detailed investigation of the term “bias”.

### 3.1 Lexical Biases (TOXTRIG)

Current toxic language detection systems often rely on the presence or absence of certain words (e.g., swear words, identity mentions) to make their predictions (Dixon et al., 2018; Dinan et al., 2019). While most previous analyses of this bias relied on a simple list of “bad” words (Davidson et al., 2019; Dinan et al., 2019),<sup>2</sup> we take a more nuanced view of how lexical items can convey toxicity, inspired by work in pragmatics and sociolinguistics of rudeness (Dynel, 2015; Kasper, 1990, *inter alia*). Specifically, we manually split our full list of words into three distinct categories depending on the extent to which they carry profane or hateful meanings or are simply associated with hateful contexts.<sup>3</sup> We refer to the full set of words as TOXTRIG, for Toxicity Triggers, which is included in our released repository.<sup>4</sup>

**Non-offensive minority identity mentions (NOI)** refers to descriptive mentions of minoritized demographic or social identities (e.g., *gay, female, Muslim*). While these mentions are not usually inherently offensive by themselves, they are often found in offensive statements that are hateful towards minorities (Dixon et al., 2018). We detect these identity mentions in text using a list of 26 regular expressions which are publicly available along with our released code.

**Possibly offensive minority identity mentions (OI)** are mentions of minoritized identities that could denote profanity or hate depending on pragmatic and contextual interpretations. This includes slurs and objectifying outdated terms to refer to minority groups, which are usually understood as attacks. Additionally, this includes *reclaimed* slurs (*queer, n\*gga*), which connote less offensive intent when spoken by in-group members compared to out-group members (Croom, 2013).

---

<sup>2</sup><https://tinyurl.com/list-of-bad-words>

<sup>3</sup>We note, however, that this categorization is itself subjective.

<sup>4</sup>[https://github.com/XuhuiZhou/Toxic\\_Debias/blob/master/data/word\\_based\\_bias\\_list.csv](https://github.com/XuhuiZhou/Toxic_Debias/blob/master/data/word_based_bias_list.csv)

**Possibly offensive non-identity mentions (ONI)** contains swear words and other profanities, which are usually offensive but not associated to any social groups (e.g., *f\*ck*, *sh\*t*). Note that the pragmatic interpretation of these words is not necessarily always toxic or offensive (Dyner, 2012), as they are often used to convey closeness between the speaker and listener or emphasize the emotionality of a statement (e.g., second example in in Figure 1.1).

### 3.2 *Dialectal Biases (AAE)*

Current toxic language detection systems also associate higher toxicity with dialectal markers of African American English (AAE) (Sap et al., 2019; Davidson et al., 2019). Since AAE is a variety of English that is common among African Americans and often signals a cultural identity in the US (Green, 2002), this dialect-based racial bias causes speech by Black authors to be suppressed more often than non-Black authors (Sap et al., 2019), thereby exacerbating racial inequality (Rosa, 2019).

In our experiments, we estimate the dialect of a tweet using a topic model from Blodgett et al. (2016). This model was trained on 60M tweets, where the dialect of the tweet was inferred from the model coordinates, which yielded a probability of a tweet being in one of four dialects (African-American English, white-aligned English, Hispanic, and other). In this study, we only focus on African-American English (AAE) and white-aligned English (WAE) tweets; both definitions are based on US English, as per Blodgett et al. (2016).<sup>5</sup> Our experiments either use the probability of a tweet being in these dialects, or assign tweets their estimated-most-probable dialect.

### 3.3 *Dataset for Toxic Language Detection*

We focus our analyses on a widely used hate speech dataset of English tweets (Founta et al., 2018). The tweets were collected using a multi-round bootstrapping procedure, and were labeled out of

---

<sup>5</sup>We avoid using disputed terms such as *general American English*, *standard American English*, or *mainstream US English*, which are frequently used for WAE, since we believe that no dialect should be privileged with the designation “general”, “standard”, or “mainstream” (Rosa, 2019).

context<sup>6</sup> for toxic language. We focus on the 86k tweets that are annotated as hateful, abusive, or neither and discard those labelled as spam. We aggregate the abusive and hateful labels into a single *toxic* category, yielding 32k toxic and 54k non-toxic tweets.<sup>7</sup>

---

<sup>6</sup>Only the tweet text—no profile information or conversational context—was shown to annotators.

<sup>7</sup>We also explored using another widely used hate speech dataset (Davidson et al., 2017), which collected tweets using a seed list of swear words and slurs. However, in line with findings by Xia et al. (2020), debiasing led to degenerate behavior due to the data collection process, as discussed in Appendix B.

## Chapter 4

### DEBIASING METHODS

We consider two types of debiasing methods from current literature. The first type addresses known, pre-defined biases—such as lexical and dialectal biases for hate speech detection, via a model-based approach involving additional training objectives (§4.1). In contrast, the second type is agnostic to prior knowledge about biases, and instead filters out examples that appear “too easy” and might hence contain spurious correlations (§4.2).

#### **4.1 *Debiased Training for Pre-Defined Toxicity Biases***

We use the LEARNED-MIXIN method of Clark et al. (2019), which achieved high out-of-distribution (OOD) performance on several NLU tasks, for debiased training. This method trains an ensemble containing a *bias-only* model which only uses pre-defined features corresponding to known biases, and a *full* model which uses all features. Intuitively, the ensemble encourages the full model to rely more on features unrelated to the biases. Once trained, the bias-only model is discarded, and only the “bias-free” full model is used for inference, following Clark et al. (2019).

**Bias-only model** Given its effectiveness on bag-of-words (BoW) features, we use an SVM classifier as the lexical-bias-only model. For example, the TOXTRIG-only model counts the frequency of TOXTRIG words in each tweet. Our dialectal-bias-only model uses the probability of dialects (AAE, WAE, Hispanic, and other) obtained from a dialect detector (Blodgett et al., 2016) as features in a SVM classifier.

**Full model** We fine-tune a RoBERTa-large classifier (Liu et al., 2019), a state-of-the-art classifier for the toxicity detection task. See Appendix A.1 for more modeling details.

Note that we only consider the `LEARNED-MIXIN-ONI` and `LEARNED-MIXIN-TOXTRIG` models for lexical debiasing, due to poor accuracies of the bias-only models for NOI and OI.<sup>1</sup>

## 4.2 Data Filtering for Spurious Biases

In addition to debiasing methods that handle known biases, we also explore automated approaches which filter out instances exhibiting unspecified, spurious biases. Specifically, we describe below two data selection methods that have shown strong OOD performance.

**AFLite (Bras et al., 2020)** is an algorithm based on the key intuition that examples predicted correctly by the simplest methods likely exhibit spurious biases. An ensemble of simple linear classifiers is trained and tested on different partitions of the data; test instances which are “predictable”, or classified correctly by most classifiers in the ensemble are discarded. The algorithm is iterative, and is repeated until a target data size is achieved. Models trained on this filtered dataset achieve higher performance on OOD and adversarially constructed test sets, compared to the original model, on several text and image classification datasets. This indicates a reduction in spurious biases in the filtered data.

**DataMaps (Swayamdipta et al., 2020)** show the presence of distinct regions in a dataset—namely, easy, hard and ambiguous—defined with respect to a given model. These regions are discovered based on the training dynamics of a model, determined by the model’s confidence in the true class, for each example, as well as the variability of this confidence, throughout training epochs. Swayamdipta et al. (2020) show that training exclusively on the hard and ambiguous regions of the data results in high OOD performance, indicating lower prevalence of spurious biases. The easy region is the largest in size for RoBERTa; however, experiments showed that training exclusively on these examples hurt OOD generalization on different NLU tasks. Following

---

<sup>1</sup>The NOI and OI bias-only models reach 63% and 67% accuracy, respectively, which is empirically hard for the ensemble to use. This is likely due to low coverage in the train set of those categories (4.43% NOI and 4.25% OI).

this work, we create DataMaps-Easy, DataMaps-Ambiguous, and DataMaps-Hard subsets for our dataset.

Following Swayamdipta et al. (2020), we set the target filtered subset size to 33% of the original training set for both filtering methods, but our filtering additionally preserved the original label proportions. We then fine-tune a RoBERTa-large classifier on these filtered subsets; see Appendix A.2 for more details.

## Chapter 5

### EXPERIMENTS: LEXICAL BIASES

We investigate the effect of debiasing approaches (§4) on removing lexical biases in hate speech detection. First, we discuss the evaluation framework for measuring bias reduction (§5.1). We present quantitative (§5.2) and qualitative (§5.3) results on lexical bias removal for all debiasing approaches, and OOD evaluation for debiased training methods (§5.4). See Appendix A.3 for hyperparameters and other experimental settings.

#### 5.1 Evaluation Framework

We report the performance of all models as overall accuracy and  $F_1$  with respect to the toxic class. Given that current hate speech systems tend to rely heavily on the presence of NOI, OI, and ONI mentions (§3.1) for labeling text as toxic, we use false positive rate (FPR) over each of these categories to measure the degree of bias in the model, following Hardt et al. (2016) and Xia et al. (2020). Specifically, we report the FPR of a model on tweets containing NOI ( $FPR_{NOI}$ ), OI ( $FPR_{OI}$ ), and ONI ( $FPR_{ONI}$ ), as well the  $F_1$  corresponding to each of these classes. Intuitively, the lower the  $FPR_*$ , the less the model infers lexical associations for toxicity, and hence is less biased.

**Evaluation for Filtered Datasets** We additionally consider metrics based on spurious lexical associations for data filtering approaches. This measures prevalence of spurious surface patterns in the filtered datasets, which might propagate to models trained on the data. Specifically, we report the Pearson’s correlation between the gold standard toxicity label and whether or not it contains NOI, OI, or ONI mentions. These correlations are denoted as  $R_{ONI}$ ,  $R_{NOI}$ , and  $R_{OI}$ , respectively; lower values indicate reduction in lexical biases.

	$R_{\text{NOI}} \downarrow$	$R_{\text{OI}} \downarrow$	$R_{\text{ONI}} \downarrow$	
Original	0.0445	0.2641	0.6718	
Random	0.0345	0.2603	0.6683	
33% train	AFLite	0.0434	0.2458	0.6016
	DataMaps-Ambig.	0.0126	0.1968	<b>0.5839</b>
	DataMaps-Hard	<b>0.0081</b>	<b>0.1853</b>	0.5849
	DataMaps-Easy	0.0772	0.3661	0.7720

Table 5.1: Lexical associations between toxicity and TOXTRIG mentions in the original dataset (Founta et al., 2018) and various filtered counterparts. Random, AFLite, and DataMaps all contain only 33% of the original data after filtering. Lower Pearson  $R$  correlation value indicates less superficial patterns in the dataset, i.e., less bias. **Takeaway:** The hard and ambiguous subsets given by DataMaps contain the lowest amount of lexical associations, indicated in boldface.

**Baselines** We consider comparison against two natural baselines: a vanilla RoBERTa-large classifier trained on the original dataset (Original). We also consider a baseline trained on a random selection of the training data (Random), for comparison with data filtering methods for debiasing. Each subset is trained on 33% of the training data.

## 5.2 Results for Lexical Bias Reduction

First, we measure the reduction in lexical biases in filtered datasets, as given by AFLite and DataMaps. As shown in Table 5.1, subsets given by AFLite and the ambiguous and hard regions produced by DataMaps reduce the overall associations between TOXTRIG words and toxicity, compared to the original and random baselines; DataMaps-Hard has the largest reduction (82% reduction). On the other hand, as expected, DataMaps-Easy shows an *increased* association between TOXTRIG mentions and toxicity, showing that these examples display overt lexical biases.

Table 5.2 shows results for lexical bias reduction using both debiased training approaches, as

	Test (12893)		NOI (602)		OI (553)		ONI (3236)		
	Acc.↑	$F_1$ ↑	$F_1$ ↑	FPR <sub>NOI</sub> ↓	$F_1$ ↑	FPR <sub>OI</sub> ↓	$F_1$ ↑	FPR <sub>ONI</sub> ↓	
Vanilla	94.21 <sub>0.0</sub>	92.33 <sub>0.0</sub>	89.76 <sub>0.3</sub>	10.24 <sub>1.3</sub>	98.84 <sub>0.1</sub>	85.71 <sub>0.0</sub>	97.34 <sub>0.1</sub>	64.72 <sub>0.8</sub>	
LMIXIN-ONI	89.65 <sub>1.5</sub>	85.59 <sub>2.5</sub>	87.04 <sub>1.1</sub>	13.99 <sub>1.5</sub>	98.87 <sub>0.0</sub>	85.71 <sub>0.0</sub>	87.87 <sub>4.5</sub>	<b>43.74</b> <sub>3.1</sub>	
LMIXIN-TOXTRIG	90.44 <sub>0.7</sub>	86.94 <sub>1.1</sub>	85.47 <sub>0.3</sub>	11.15 <sub>1.7</sub>	97.64 <sub>0.3</sub>	<b>71.43</b> <sub>0.0</sub>	90.41 <sub>1.8</sub>	44.55 <sub>1.5</sub>	
33% train	Random	94.07 <sub>0.1</sub>	92.18 <sub>0.1</sub>	89.48 <sub>0.4</sub>	9.33 <sub>0.7</sub>	<b>98.93</b> <sub>0.0</sub>	<b>83.33</b> <sub>3.4</sub>	97.40 <sub>0.1</sub>	67.15 <sub>0.6</sub>
	AFLite	93.86 <sub>0.1</sub>	91.94 <sub>0.1</sub>	<b>90.21</b> <sub>0.4</sub>	11.26 <sub>1.1</sub>	98.90 <sub>0.0</sub>	85.71 <sub>0.0</sub>	97.32 <sub>0.1</sub>	67.97 <sub>3.4</sub>
	DataMaps-Ambig.	94.33 <sub>0.1</sub>	92.45 <sub>0.1</sub>	89.16 <sub>0.7</sub>	7.39 <sub>1.0</sub>	98.87 <sub>0.0</sub>	85.71 <sub>0.0</sub>	<b>97.54</b> <sub>0.0</sub>	64.39 <sub>1.4</sub>
	DataMaps-Hard	<b>94.50</b> <sub>0.0</sub>	<b>92.61</b> <sub>0.1</sub>	89.54 <sub>0.4</sub>	6.26 <sub>0.9</sub>	98.84 <sub>0.0</sub>	85.71 <sub>0.0</sub>	97.43 <sub>0.0</sub>	61.95 <sub>1.1</sub>
	DataMaps-Easy	94.00 <sub>0.1</sub>	91.94 <sub>0.2</sub>	86.81 <sub>0.6</sub>	<b>5.92</b> <sub>0.7</sub>	98.87 <sub>0.0</sub>	<b>83.33</b> <sub>3.4</sub>	97.17 <sub>0.1</sub>	<b>60.33</b> <sub>3.8</sub>

Table 5.2: Evaluation of lexical bias removal for all debiasing methods on the Founta et al. (2018) test set. Results show the mean and s.d. (subscript) of accuracy and  $F_1$  across 3 runs, as well as  $F_1$  and false positive rate exclusively on test examples containing specific TOXTRIG mentions—NOI, OI and ONI, along with the number of examples in each category. The lower the FPR\*, the less the model infers lexical associations for toxicity. The first block shows debiased training approaches, along with the vanilla classifier, which are trained on the full dataset. The second block shows data filtering approaches, all trained on only 33% of the training data. Best performance in each block is boldfaced. *Takeaway*: While data filtering approaches achieve overall higher performance, debiased training approaches perform better on lexical bias reduction, in aggregate.

well as models trained on datasets filtered using AFLite and all three regions from DataMaps. Both debiased training approaches, LMIXIN-ONI and LMIXIN-TOXTRIG, reduce FPR<sub>ONI</sub> as well as FPR<sub>OI</sub> by a large amount. However, both approaches also hurt in-distribution test performance, indicating that ONI and other TOXTRIG features are essential for good performance.<sup>1</sup> In contrast, the models trained on hard and ambiguous subsets from DataMaps both preserve in-distribution performance, even though they are trained only a third of the original data. They also reduce the

<sup>1</sup>When we combine the bias-only model and the full model, we obtain competitive performance (see Appendix A.4).

	Gold	DM- Hard	DM- Easy
@user THIS? LMAOOOOO...do yourself a solid and stay out of Black people’s men- tions and mind your caucasia...	🍃	💀	🍃
RT @user I wish I wasn’t so annoying like I even piss myself off	💀	🍃	💀
@user If you want to attack people, attack fundamentalists of all faiths. Attack those who condemn 1.5bn people out of hand.	🍃	💀	🍃

Table 5.3: Examples of test set tweets with their gold-standard annotations and predictions from models trained on DataMaps-Hard (DM-Hard) and DataMaps-Easy (DM-Easy) subsets. 💀 denotes tweets with toxic labels, and 🍃 represents non-toxic labels. We anonymize the usernames to protect user privacy.

rate of falsely predicting NOI mentions as toxic ( $FPR_{NOI}$ ), while not showing much improvement for ONI and maintaining  $FPR_{OI}$  of the original baseline.

Surprisingly, the model trained on the easy subset from DataMaps shows good bias reduction on the NOI and ONI categories, while matching the random selection baseline for OI. This is despite DataMaps-Easy showing an increased association between TOXTRIG mentions and toxicity (Table 5.1). Notably, the  $F_1$  for all categories suffers under this model, indicating that it is less competent than the baseline. These results suggest that reduced associations in the data might not necessarily lead to debiased models trained on the same data. Overall, no single approach outperforms all others across different categories for lexical debiasing.

### 5.3 Qualitative Analysis

A qualitative study of the Founta et al. (2018) test set shows the presence of many annotation errors. We show three representative annotation errors in Table 5.3. The first example contains an atypical example of toxicity, towards white folks, which the annotators might have been unaware of. It also contains a link which annotators had access to, but not models. The second contains

the word *p\*ss* which the annotators may have relied on for their assessment. The third encourages violence/abuse towards an identity which isn't typically the target of violence. Interestingly, the DataMaps-Easy predictions agree with all the gold standard annotations; perhaps such annotation errors and ambiguity are responsible for the performance discussed in §5.2. These annotation ambiguities might also impair our measurement for models' performance and debiasing effects, and expose a limitation of these datasets.

#### **5.4 Adversarial Evaluation: ONI-Adv**

To further study the reliance of debiased models on the ONI words, we use the test set from Dinan et al. (2019), which contains 3000 manually created sentences, 300 toxic. The toxic sentences barely contain any ONI mentions, e.g., *Cheese is made by french people who smell*. We call this test set ONI-Adv (for adversarial) since it challenges models with a reversal in the association between toxicity and offensive non-identity words (e.g., "*f\*ck*", "*sh\*t*").

We report  $F_1$  for all models in Figure 5.1, which shows how well a model identifies toxicity in offensive tweets that do not contain overtly lexical cues of toxicity. The debiased training approaches improve over the baselines; data filtering methods do not. One reason for this might be that data filtering methods were trained on much less data than both LMIXIN models. Regardless, none of the models we test are good at predicting subtle, non-overt toxicity.

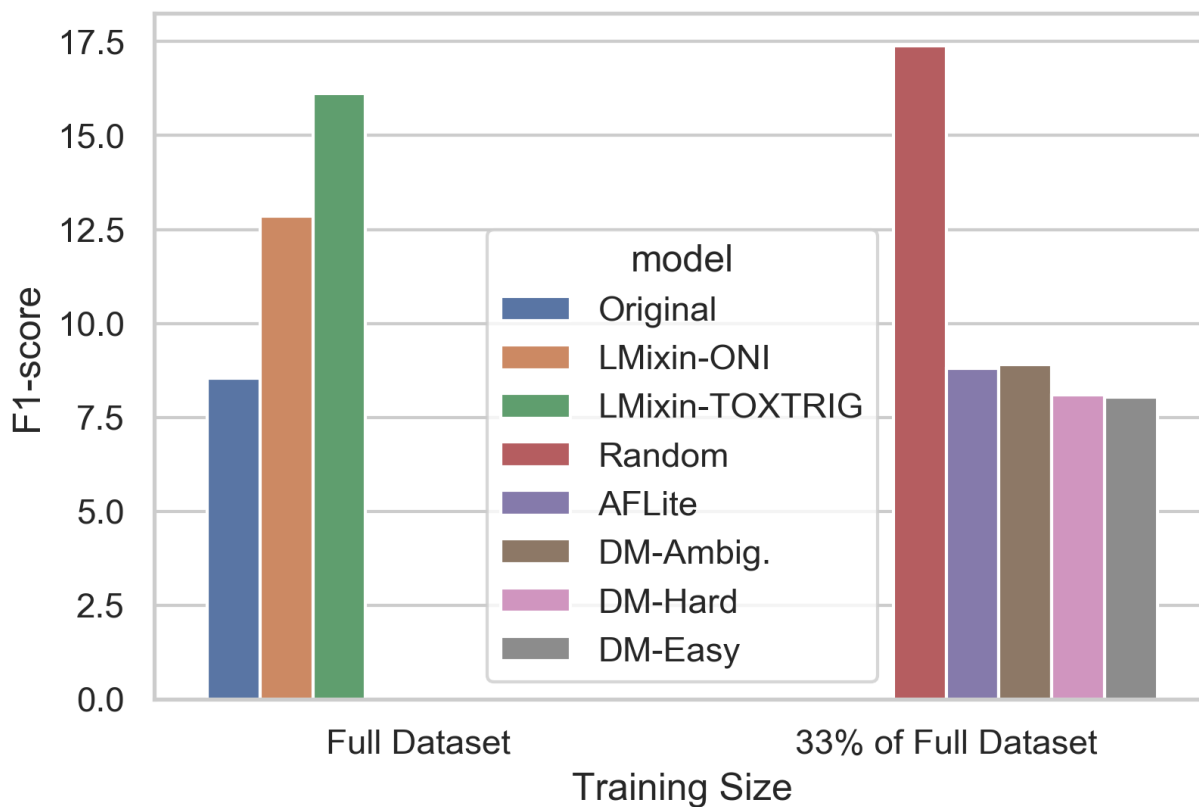


Figure 5.1: Challenge set evaluation for lexical biases, comparing all debiasing methods with baselines, using the ONI-Adv test set. **Takeaway:**  $F_1(\uparrow)$  measures show that all models perform poorly at identifying toxic text not containing overtly lexical cues of toxicity. In general, debiased training approaches outperform the original model on this challenge set, while data filtering is not as effective.

## Chapter 6

### EXPERIMENTS: DIALECTAL AND RACIAL BIASES

We test the efficacy of the bias reduction methods from §4 for dialectal bias (§3.2) reduction.

#### 6.1 *Dialectal Biases*

For our dialectal bias experiments, we first infer the dialect of a tweet as described in §3.2. Then, analogous to the lexical bias evaluation, we quantify the dialectal debiasing using the Pearson’s correlation between estimated probabilities of AAE and toxicity ( $R_{AAE}$ ), and the false positive rates of models on AAE tweets ( $FPR_{AAE}$ ). See Appendix A.3 for hyperparameters and other experimental settings.

Results in Table 6.1 show that almost all data filtering and debiasing methods reduce dialectal biases, with DataMaps-Easy as the exception (consistent with Table 5.1). Notably, DataMaps-Hard performs the best at dialectal debiasing, both in terms of toxicity-AAE correlation ( $R_{AAE}$ ) and in terms of false flagging of toxicity ( $FPR_{AAE}$ ). Interestingly, most models’ decrease in false flagging is small, suggesting room for improvement.

#### 6.2 *Racial Biases*

To quantify the real-world impact of dialect-based racial bias, we measure the rates of toxicity predicted by models on a corpus of tweets for which the race of authors is available, but not annotations of toxicity. Specifically, we consider the dataset released by Preoțiuc-Pietro and Ungar (2018), which consists of 5.4M tweets, collected from 4,132 survey participants (3,184 White, 374 African American) with self-reported race/ethnicity and Twitter user handles.<sup>1</sup>

---

<sup>1</sup>For efficiency, we randomly select 12k tweets from the dataset as the OOD test set.

We quantify our models' racial bias by measuring the difference in rates of flagging tweets by African American authors and those by white authors, following Sap et al. (2019).<sup>2</sup>

Listed in Table 6.2, our results show that automatic debiasing methods do not consistently decrease the racial discrepancy in flagging toxicity. Notably, the toxicity rates on tweets by African American authors—and the differences compared to white authors—are similar across all debiasing methods and baselines, except for DataMaps-Easy, which shows the most racial bias in toxicity flagging. Surprisingly, DataMaps-Hard, which mitigated dialectal bias the best out of all debiasing methods, also shows high discrepancy between author races. Confirming previous results, this suggests that debiasing these systems requires more than automatic debiasing methods.

---

<sup>2</sup>Note that we assume that authors from all races have the same likelihood of writing toxic language.

		Test		
		$R_{AAE} \downarrow$	$F_1 \uparrow$	$FPR_{AAE} \downarrow$
Vanilla		0.4079	92.33 <sub>0.0</sub>	16.84 <sub>0.3</sub>
LMIXIN-Dialect		-	92.26 <sub>0.1</sub>	16.07 <sub>0.4</sub>
33% train	Random	0.4027	92.18 <sub>0.1</sub>	16.67 <sub>0.6</sub>
	AFLite	0.3577	91.94 <sub>0.1</sub>	16.84 <sub>0.8</sub>
	DataMaps-Ambig.	0.2965	92.45 <sub>0.1</sub>	15.99 <sub>0.4</sub>
	DataMaps-Hard	<b>0.2878</b>	<b>92.61</b> <sub>0.1</sub>	<b>13.71</b> <sub>0.2</sub>
	DataMaps-Easy	0.5347	91.94 <sub>0.2</sub>	19.46 <sub>2.8</sub>
AAE-relabeled		0.3453	91.64 <sub>0.3</sub>	<b>12.69</b> <sub>0.0</sub>

Table 6.1: Dialectal bias evaluation for all debiasing methods (§6), as well as the relabeling approach (§7) on the Founta et al. (2018) test set. We report  $F_1$  and the false positive rate with respect to tweets in AAE ( $FPR_{AAE}$ ), reflecting dialectal bias (lower is less biased), showing mean and s.d. (subscript) across 3 runs. (Top Block) Debaised training approaches, along with the vanilla classifier, are all trained on the full dataset. (Middle Block) Random, AFLite and DataMaps all are trained on only 33% of the training data. Best performance for each training set size is in boldface. **Takeaway:** Both debiasing approaches improve performance over baselines, with DataMaps-Hard proving the most effective at debiasing. (Bottom Block) AAE-relabeling results in a model which despite following a noisy process yields even larger improvements for dialectal debiasing.

	W-Tox.	AA-Tox.	$\Delta$	$\downarrow$ AA/W $\downarrow$
Original	7.24	12.61	5.37	1.74
LMIXIN-Dialect	7.50	12.55	5.06	1.67
Random	8.28	13.24	4.96	1.60
AFLite	7.32	11.64	4.33	1.59
DataMaps-Ambig.	6.75	12.17	5.42	1.80
DataMaps-Hard	6.36	11.67	5.31	1.84
DataMaps-Easy	8.46	16.30	7.83	1.94
AAE-relabeled	6.93	10.60	<b>3.67</b>	<b>1.53</b>

Table 6.2: Racial disparity in toxicity prediction reported on Preoțiu-Pietro and Ungar (2018). **W-Tox.** indicates % of white users’ tweets being flagged as toxic, **AA-Tox.** indicates % of African American users’ tweets being flagged as toxic,  $\Delta$  refers to the difference between AA-Tox. and W-Tox., and **AA/W** refers to the ratio between AA-Tox. and W-Tox. **Takeaway:** Methods generally fail in debiasing on this OOD test set except the relabeling approach shows some benefit.

## Chapter 7

### TOWARDS DATA RELABELING

Based on our quantitative and qualitative analyses, we believe there still is room for improvement in debiasing hate speech detection. Therefore, we turn our attention to the role of label noise in datasets. Partly inspired by our qualitative analyses of debiased models’ predictions, we design a proof-of-concept study where we automatically correct the label of tweets using a(n automatic) dialectal translation of the tweet, inspired by previous work showing that highlighting AAE tweets’ dialect led them to be labeled as less toxic (Sap et al., 2019). We conclude this study by discussing the limitations and ethical implications of the synthetic data, and cautioning against its real-world application.

Focusing on dialectal bias, our key assumption is that an AAE tweet and its corresponding WAE version should have the same toxicity label, therefore toxic AAE tweets whose WAE versions are non-toxic are candidates for label correction.<sup>1</sup>

However, gold-standard translations of AAE to WAE would require qualified translators, and automatic AAE-to-WAE translation systems do not exist, to the best of our knowledge. Therefore, we create a proof-of-concept study—we set up a AAE to WAE “translation” system using the few-shot capabilities of the GPT-3 language model (Brown et al., 2020). Under this mechanism, we prompt GPT-3 with four translation pairs (taken from (Spears, 1998)) and an AAE tweet from our training data, and generate its WAE “translation”. The list of prompts, as well as further details, are provided in Appendix C. Note that we do *not* recommend this approach to build large scale parallel data for dialects, as discussed under ethical implications and limitations.

Next, as per our heuristic, we only relabel toxic AAE tweets whose WAE translation is predicted

---

<sup>1</sup>Note that this assumption does not hold for lexical items, because substituting lexical items (e.g., swapping a minority mention for a majority mention) would drastically change the denotational meaning of the sentence.

AAE	GPT-3 WAE Translation	Gold	New
RT @user I can't stand a bad texter bruh like don't be mad if I forget about yo ass	RT @user I can't stand a bad texter bro like don't be mad if I forget about you	☠	🍃
RT @user Retweet if you fuck with this!!!!	RT @user Retweet if you like this!	☠	🍃
RT @user That nigga needs anger management	RT @user That guy needs anger management	☠	🍃
RT @user oh fucking hell take a day off man	RT @user oh fuck take a day off man	☠	☠

Table 7.1: Examples of AAE tweets with their GPT-3 based WAE translation, and original gold standard and new annotations based on AAE-relabeled. For the first three tweets, the (biased) gold labels are changed by models predicting the new labels on their WAE translations. ☠ indicates presence of toxicity, and 🍃 represents non-toxic. We anonymize the usernames to protect user privacy.

as non-toxic by either our vanilla classifier trained on the original Founta et al. (2018) dataset, or an identical classifier trained on the WAE translated tweets. The resulting dataset (AAE-relabeled) is the same size as the original dataset, but with 954 (12%) out of 8260 toxic AAE tweets relabeled as non-toxic (examples in Table 7.1). To assess the validity of the relabeling, the first three authors manually annotated toxicity of 50 randomly selected relabeled tweets. On average, authors agreed with 84% of the relabeling decisions.

Then, we evaluate the dialectal bias of AAE-relabeled and quantify the dialect and racial prediction biases from a RoBERTa-large classifier trained on AAE-relabeled, following §6. As shown in the last row of Table 6.1, this relabeling scheme decreases dialectal bias more than any other debiasing method, specifically as measured by the FPR on AAE tweets, with one point drop in  $F_1$  score. The  $F_1$  score on the “gold” test data (Table 6.1) are not fully reliable, as test data contain label biases and better performance could come from exploiting these biases. As shown in Table 6.2, the model trained on AAE-relabeled has the lowest racial disparity in toxicity flagging rates compared to all other methods.

These results highlight that debiasing methods are much less effective at mitigating dialectal dataset biases compared to data relabeling. For future investigations, we recommend obtaining human-written AAE-WAE pairs (e.g., as done by (Groenwold et al., 2020). ) Additionally, to ensure less biased toxicity labeling, we recommend recruiting AAE speakers or experts for avoiding over-identification of AAE-markers as toxic (Spears, 1998; Croom, 2013). Alternatively, we recommend exploring more holistic representations of social biases or toxicity (e.g., Social Bias Frames) (Sap et al., 2020).

### ***Ethical Implications & Limitations***

The above synthetic setting is meant to illustrate the role of labeling quality on biases in annotations. We strongly caution against using this approach in real-world applications, such as building parallel datasets for dialects. First, due to how its training data was selected, GPT-3 has likely not been exposed to many African American English varieties during training (Jo and Gebru, 2020). Second, pretrained language models are known to generate toxic language at non-trivial rates (Gehman et al., 2020), which could cause differential toxicity in the translations.

## Chapter 8

### **CONCLUSION**

We investigate whether toxic language detection systems can be debiased using recently introduced methods for debiasing text classification in NLU tasks. Focusing on two types of biases, lexical and dialectal, our experiments show that these methods face significant challenges in reducing the biased behavior in toxicity detectors. This indicates that biases in toxic language detection might be different in nature compared to spurious associations studied in typical NLU settings. We studied a synthetic scheme for relabeling examples with potential dialectal biases; our results indicate that correcting noisy labels results in better bias reduction. Our findings suggest that instead of solely relying on development of automatic debiasing for existing, imperfect datasets, future work should focus primarily on the quality of the underlying data for hate speech detection, such as accounting for speaker identity and dialect. Indeed, such efforts could act as an important step towards making systems less discriminatory, and hence safe and usable.

## BIBLIOGRAPHY

Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Proc. of WSM*, 2018. URL <https://arxiv.org/abs/1802.00393>.

Daniel Preoțiuc-Pietro and Lyle Ungar. User-level race and ethnicity predictors from twitter text. In *Proc. of COLING*, 2018. URL <https://www.aclweb.org/anthology/C18-1130>.

Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, 2017. URL <https://ojs.aaai.org/index.php/ICWSM/article/view/14955>.

Danyaal Yasin. Black and banned: Who is free speech for?, 2018. URL <https://www.indexoncensorship.org/2018/09/black-and-banned-who-is-free-speech-for/>.

Jessica Guynn. What civil rights groups want from facebook boycott: Stop hate speech and harassment of black users, 2020. URL <https://web.archive.org/save/https://www.usatoday.com/story/tech/2020/07/07/facebook-ad-boycott-racism-harassment-hate-african-americans/5385514002/>.

Jae Yeon Kim, Carlos Ortiz, Sarah Nam, Sarah Santiago, and Vivek Datta. Intersectional bias in hate speech and abusive language datasets, 2020.

Thiago Dias Oliva, Dennys Marcelo Antonialli, and Alessandra Gomes. Fighting hate

- speech, silencing drag queens? artificial intelligence in content moderation and risks to lgbtq voices online. In *Sexuality & Culture*, 2020. URL <https://doi.org/10.1007/s12119-020-09790-w>.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. The risk of racial bias in hate speech detection. In *Proc. of ACL*, 2019. URL <https://www.aclweb.org/anthology/P19-1163>.
- Lucas Dixon, John Li, Jeffrey Scott Sorensen, Nithum Thain, and L. Vasserman. Measuring and mitigating unintended bias in text classification. In *Proc. of AES*, 2018. URL <https://dl.acm.org/doi/pdf/10.1145/3278721.3278729>.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. Build it break it fix it for dialogue safety: Robustness from adversarial human attack. In *Proc. of EMNLP*, 2019. URL <https://www.aclweb.org/anthology/D19-1461>.
- Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. Racial bias in hate speech and abusive language detection datasets. In *Abusive Language Workshop (at ACL)*, 2019. URL <https://www.aclweb.org/anthology/W19-3504>.
- Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proc. of AES*. Association for Computing Machinery, 2018. URL <https://doi.org/10.1145/3278721.3278779>.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proc. of EMNLP*, 2019. URL <https://www.aclweb.org/anthology/D19-1418>.
- He He, Sheng Zha, and Haohan Wang. Unlearn dataset bias in natural language inference by fitting the residual. In *EMNLP Workshop on Deep Learning Approaches for Low-Resource NLP*, 2019. URL <https://www.aclweb.org/anthology/D19-6115>.

Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. End-to-end bias mitigation by modelling biases in corpora. In *Proc. of ACL, 2020*. URL <https://www.aclweb.org/anthology/2020.acl-main.769>.

Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. Adversarial filters of dataset biases. In *Proc. of ICML, 2020*. URL <https://arxiv.org/abs/2002.04108>.

Mengzhou Xia, Anjalie Field, and Yulia Tsvetkov. Demoting racial bias in hate speech detection. In *Proc. of Social NLP, 2020*. URL <https://www.aclweb.org/anthology/2020.socialnlp-1.2>.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Proc. of NeurIPS, 2020*. URL <https://papers.nips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.

Ellery Wulczyn, Nithum Thain, and Lucas Dixon. Ex machina: Personal attacks seen at scale. In *Proc. of WWW. International World Wide Web Conferences Steering Committee, 2017*. URL <https://doi.org/10.1145/3038912.3052591>.

David Noever. Machine learning suites for online toxicity detection. In *Proc. of CoRR, 2018*. URL <http://arxiv.org/abs/1810.01869>.

Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. Social bias frames: Reasoning about social and power implications of language. In *Proc. of ACL, 2020*. URL <https://www.aclweb.org/anthology/2020.acl-main.486>.

Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Athena Vakali. Mean birds: Detecting aggression and bullying on twitter. In *Proc. of ACM WSC*. ACM, 2017. URL <http://doi.acm.org/10.1145/3091478.3091487>.

Jennifer Golbeck, Zahra Ashktorab, Rashad O. Banjo, Alexandra Berlinger, Siddharth Bhagwan, Cody Buntain, Paul Cheakalos, Alicia A. Geller, Quint Gergory, Rajesh Kumar Gnanasekaran, Raja Rajan Gunasekaran, Kelly M. Hoffman, Jenny Hottle, Vichita Jienjiltert, Shivika Khare, Ryan Lau, Marianna J. Martindale, Shalmali Naik, Heather L. Nixon, Piyush Ramachandran, Kristine M. Rogers, Lisa Rogers, Meghna Sardana Sarin, Gaurav Shahane, Jayanee Thanki, Priyanka Vengataraman, Zijian Wan, and Derek Michael Wu. A large labeled corpus for online harassment research. In *Proc. of Web Science Conference*. ACM, 2017. URL <http://doi.acm.org/10.1145/3091478.3091509>.

Zeerak Waseem and Dirk Hovy. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proc. of NAACL SRW*, 2016. URL <http://www.aclweb.org/anthology/N16-2013>.

Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. Benchmarking aggression identification in social media. In *Proc. of TRAC*, 2018. URL <https://www.aclweb.org/anthology/W18-4401>.

Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. Deep learning for hate speech detection in tweets. In *Alan Turing Institute*, 2017.

John Pavlopoulos, Prodromos Malakasiotis, and Ion Androutsopoulos. Deep learning for user comment moderation. In *Proc. of ALW*, 2017. URL <https://www.aclweb.org/anthology/W17-3004>.

Keita Kurita, Anna Belova, and Antonios Anastasopoulos. Towards robust toxic content classification, 2019.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL*, 2019. URL <https://www.aclweb.org/anthology/N19-1423>.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- Ji Ho Park, Jamin Shin, and Pascale Fung. Reducing gender bias in abusive language detection. In *Proc. of EMNLP*, 2018. URL <https://www.aclweb.org/anthology/D18-1302>.
- Ameya Vaidya, Feng Mai, and Yue Ning. Empirical analysis of multi-task learning for reducing model bias in toxic comment detection. In *Proc. of ICWSM*, 2019. URL <https://arxiv.org/abs/1909.09758>.
- Marzieh Mozafari, Reza Farahbakhsh, and Noël Crespi. Hate speech detection and racial bias mitigation in social media based on bert model. In *PLOS ONE*. Public Library of Science, 2020. URL <https://doi.org/10.1371/journal.pone.0237861>.
- Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and V. Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proc. of ICCV*, 2019. URL <https://arxiv.org/abs/1811.08489>.
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. Mind the trade-off: Debiasing NLU models without degrading the in-distribution performance. In *Proc. of ACL*, 2020. URL <https://www.aclweb.org/anthology/2020.acl-main.770>.
- Su Lin Blodgett, Solon Barocas, Hal Daumé, III, and Hanna Wallach. Language (technology) is power: A critical survey of “bias” in NLP. In *Proc. of ACL*, 2020. URL <https://www.aclweb.org/anthology/2020.acl-main.485/>.
- Sarah T Roberts. *Behind the screen: Content moderation in the shadows of social media*. Yale University Press, 2019. URL <https://yalebooks.yale.edu/book/9780300235883/behind-screen>.

- Bertie Vidgen, Helen Margetts, and Alex Harris. How much online abuse is there? In *Alan Turing Institute*, 2019. URL [https://www.turing.ac.uk/sites/default/files/2019-11/online\\_abuse\\_prevalence\\_full\\_24.11.2019\\_-\\_formatted\\_0.pdf](https://www.turing.ac.uk/sites/default/files/2019-11/online_abuse_prevalence_full_24.11.2019_-_formatted_0.pdf).
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. In *Proc. of EMNLP*, 2015. URL <https://www.aclweb.org/anthology/D15-1075>.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proc. of NAACL*, 2018. URL <https://www.aclweb.org/anthology/N18-1101>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proc. of EMNLP*, pages 2383–2392, 2016. URL <https://arxiv.org/abs/1606.05250>.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wojatzki. Measuring the reliability of hate speech annotations: the case of the european refugee crisis. In *NLP 4 CMC Workshop*, 2017. URL <https://arxiv.org/abs/1701.08118>.
- Roy Schwartz, Maarten Sap, Ioannis Konstas, Li Zilles, Yejin Choi, and Noah A Smith. The effect of different writing tasks on linguistic style: A case study of the roc story cloze task. In *Proc. of CoNLL*, 2017. URL <https://www.aclweb.org/anthology/K17-1004>.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. Annotation artifacts in natural language inference data. In *Proc. of NAACL*, 2018. URL <https://arxiv.org/abs/1803.02324>.
- Arthur K Spears. African-American language use: Ideology and so-called obscenity. In *African-American English: Structure, History and Use*. Routledge New York, 1998. URL [https://www.researchgate.net/publication/253946995\\_African-American\\_Language\\_Use\\_Ideology\\_and\\_So-called\\_Obscenity](https://www.researchgate.net/publication/253946995_African-American_Language_Use_Ideology_and_So-called_Obscenity).

Björn Technau. Going beyond hate speech: The pragmatics of ethnic slur terms. *Lodz Papers in Pragmatics*, 14(1):25–43, 2018. URL <https://www.degruyter.com/view/journals/lpp/14/1/article-p25.xml?language=en>.

Jonathan Rosa and Nelson Flores. Unsettling race and language: Toward a raciolinguistic perspective. In *Language In Society*. Cambridge University Press, 2017. URL <https://www.cambridge.org/core/journals/language-in-society/article/unsettling-race-and-language-toward-a-raciolinguistic-perspective/30FFC5253F465905D75CDDFF1C1363AE3>.

Marta Dynel. The landscape of impoliteness research. In *Journal of Politeness Research*, 2015. URL [https://www.researchgate.net/publication/281927834\\_The\\_landscape\\_of\\_impoliteness\\_research](https://www.researchgate.net/publication/281927834_The_landscape_of_impoliteness_research).

Gabriele Kasper. Linguistic politeness: current research issues. In *Journal of Pragmatics*. Elsevier, 1990. URL <https://www.semanticscholar.org/paper/Linguistic-politeness%3A%3A-Current-research-issues%E2%98%86-Kasper/02009b316378e0e8472542e11225d1f602d5576f>.

Adam M Croom. How to do things with slurs: Studies in the way of derogatory words. In *Language & communication*, 2013. URL <https://philarchive.org/archive/CROHTD>.

Marta Dynel. Swearing methodologically : the (im)politeness of expletives in anonymous commentaries on youtube. In *Journal of English Studies*, 2012. URL <https://publicaciones.unirioja.es/ojs/index.php/jes/article/view/179>.

Lisa Green. *African American English: A Linguistic Introduction*. Cambridge University Press, 2002. URL <https://www.cambridge.org/core/books/african-american-english/1AE59657F9CF1BBC3A2BF2B9BB29D1D0>.

Jonathan Rosa. *Looking like a language, sounding like a race*. Oxford University Press,

2019. URL <https://oxford.universitypressscholarship.com/view/10.1093/oso/9780190634728.001.0001/oso-9780190634728>.
- Su Lin Blodgett, Lisa Green, and Brendan O’Connor. Demographic dialectal variation in social media: A case study of African-American English. In *Proc. of EMNLP*, 2016. URL <https://www.aclweb.org/anthology/D16-1120>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. In *arXiv preprint arXiv:1907.11692*, 2019. URL <https://arxiv.org/abs/1907.11692>.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proc. of EMNLP*, 2020. URL <https://arxiv.org/abs/2009.10795>.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proc. of NeurIPS*, 2016. URL <https://papers.neurips.cc/paper/6374-equality-of-opportunity-in-supervised-learning.pdf>.
- Sophie Groenwold, Lily Ou, Aesha Parekh, Samhita Honnavalli, Sharon Levy, Diba Mirza, and William Yang Wang. Investigating African-American vernacular english in Transformer-Based text generation. In *Proc. of EMNLP*, 2020. URL <https://www.aclweb.org/anthology/2020.emnlp-main.473>.
- Eun Seo Jo and Timnit Gebru. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proc. of FAT*, 2020. URL <https://dl.acm.org/doi/abs/10.1145/3351095.3372829>.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. Realtocixityprompts: Evaluating neural toxic degeneration in language models. In *Findings of EMNLP*, 2020. URL <https://arxiv.org/abs/2009.11462>.

Shilpa Arora. Active learning for natural language processing literature review. 2007.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. Huggingface’s transformers: State-of-the-art natural language processing, 2019. URL <https://arxiv.org/abs/1910.03771>.

Chris Sweeney and Maryam Najafian. A transparent framework for evaluating unintended demographic bias in word embeddings. In *Proc. of ACL*, 2019. URL <https://www.aclweb.org/anthology/P19-1162>.

Tal Linzen. How can we accelerate progress towards human-like linguistic generalization? In *Proc. of ACL*, 2020. URL <https://www.aclweb.org/anthology/2020.acl-main.465>.

Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proc. of ACL*, 2017. URL <https://www.aclweb.org/anthology/P17-1147>.

Betty van Aken, Julian Risch, Ralf Krestel, and Alexander Löser. Challenges for toxic comment classification: An in-depth error analysis. In *Proc. of ALW*, 2018.

Wei Emma Zhang, Quan Z. Sheng, and Ahoud Abdulrahmn F. Alhazmi. Generating textual adversarial examples for deep learning models: A survey. In *Proc. of CoRR*, 2019. URL <http://arxiv.org/abs/1901.06796>.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WINOGRANDE: an adversarial winograd schema challenge at scale. In *Proc. of CoRR*, 2019a. URL <http://arxiv.org/abs/1907.10641>.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WINOGRANDE: an adversarial winograd schema challenge at scale. In *Proc. of CoRR*, 2019b. URL <http://arxiv.org/abs/1907.10641>.

- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proc. of EMNLP*, 2017. URL <https://www.aclweb.org/anthology/D17-1323>.
- Yanai Elazar and Y. Goldberg. Adversarial removal of demographic attributes from text data. In *Proc. of EMNLP*, 2018. URL <https://www.aclweb.org/anthology/D18-1002/>.
- Stephen L Mann. Drag queens’ use of language and the performance of blurred gendered and racial identities. In *Journal of homosexuality*, 2011.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a large-scale dataset for fact extraction and VERification. In *Proc. of NAACL*, 2018. URL <https://www.aclweb.org/anthology/N18-1074>.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. Predicting the type and target of offensive posts in social media. In *Proc. of NAACL*, 2019. URL <https://www.aclweb.org/anthology/N19-1144>.
- Ying Chen, Y. Zhou, S. Zhu, and H. Xu. Detecting offensive language in social media to protect adolescent online safety. In *PASSAT*, 2012.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. SWAG: A large-scale adversarial dataset for grounded commonsense inference. In *Proc. of EMNLP*, 2018a. URL <https://www.aclweb.org/anthology/D18-1009>.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. SWAG: A large-scale adversarial dataset for grounded commonsense inference. In *Proc. of EMNLP*, 2018b. URL <https://www.aclweb.org/anthology/D18-1009>.
- Sachin Kumar, Shuly Wintner, Noah A. Smith, and Yulia Tsvetkov. Topics to avoid: Demoting latent confounds in text classification. In *Proc. of EMNLP*, 2019. URL <https://www.aclweb.org/anthology/D19-1425>.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proc. of ACL*, 2019. URL <https://www.aclweb.org/anthology/P19-1472>.

Tom McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proc. of ACL*.

Zeeraq Waseem. Are you a racist or am i seeing things? annotator influence on hate speech detection on twitter. In *Proc. of NLP & CSS*, 2016. URL <http://aclweb.org/anthology/W16-5618>.

Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, et al. Learning and evaluating general linguistic intelligence. In *arXiv preprint arXiv:1901.11373*, 2019.

Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sanchez, Mugdha Pandya, and Adam Lopez. Intrinsic bias metrics do not correlate with application bias. In *arXiv preprint arXiv:2012.15859*, 2020.

**APPENDIX**

## Appendix A

### FURTHER DETAILS FOR MODELS

#### A.1 Model Debiasing

The LEARNED-MIXIN ensemble allows the model to explicitly determine how much to trust the bias given the input:

$$\hat{p}_i = \text{softmax}\{\log(p_i) + g(\mathbf{x}_i) \log b_i\}$$

where  $\mathbf{x}_i$  is the  $i$ th input text,  $p_i$  and  $b_i$  is the toxicity prediction produced by RoBERTa, and bias-only model respectively, and  $g$  is a parametric function, which is defined as  $\text{softplus}(\mathbf{w} \cdot \mathbf{h}_i)$ , where  $\mathbf{w}$  is a learned vector,  $\mathbf{h}_i$  is the last hidden layer of the model for example  $\mathbf{x}_i$ , and the  $\text{softplus}(x) = \log(1 + \exp x)$ . To prevent the LEARNED-MIXIN ensemble from ignoring  $b_i$ , Clark et al. (2019) add an entropy penalty ( $H$ ) to the loss:

$$R = \alpha H(\text{softmax}\{g(\mathbf{x}_i) \log b_i\})$$

Where  $H(z) = -\sum_j z_j \log z_j$  is the entropy and  $\alpha$  is a hyperparameter.

#### A.2 Data Filtering

For the data filtering methods, we first filter data to 50% of the original data as in Swayamdipta et al. (2020). Then we further downsample the dataset to 33% of the original data to control that each training set has the same toxic ratio as the original training set. This step is to avoid confounding our results with different toxic ratio among different training sets.

#### A.3 Training Settings

For all the experiments, we fine-tune RoBERTa-large (Liu et al., 2019) over the corresponding corpus with one GTX2080 Ti. We use the default hyperparameters as provided in the HuggingFace

Transformers library (Wolf et al., 2019), with two major changes: we use a learning rate of  $10^{-5}$  and 8 batch size in all experiments.

#### ***A.4 Prediction Combining with Bias-only Model***

To prevent the possibility that our LMIXIN-TOXTRIG/ONI is not well trained, thus resulting in the decrease of models' in-distribution performance, we use the joint-prediction from the main and bias-only model to infer the in-distribution test set and they obtain 94.15% and 94.17% accuracy, respectively. This is competitive performance as shown in Table 5.2.

## Appendix B

### **ALTERNATIVE DATASET OF TOXIC LANGUAGE**

Davidson et al. (2017) collected data from Twitter, starting with 1,000 terms from HateBase (an online database of hate speech terms) as seeds, which the process relies on lexical biases. We find that performing data filtering methods over this dataset leads to degenerate behaviour. Specifically, as shown in Table B.1, the easy region demonstrates least spurious correlation due to its heavily skewed class distribution, which further prevent us from downsampling to control the toxic ratio. We also train LMIXIN-TOXTRIG and LMIXIN-dialect over the dataset. Table B.2 shows that FPR of the debiased model increase instead except for the OI category and Table B.3's results behave in-line with Table 6.1.

	Toxic Ratio	$R_{\text{NOI}} \downarrow$	$R_{\text{OI}} \downarrow$	$R_{\text{ONI}} \downarrow$	$R_{\text{AAE}} \downarrow$
Original†	0.8308	0.0287	0.4320	0.2610	0.4061
Random	0.8312	0.0288	0.4312	0.2621	0.4011
AFLite	0.7669	0.0342	0.4708	0.2835	0.4236
DataMaps-Ambig.	0.6736	0.0493	0.4683	0.3230	0.4445
DataMaps-Hard	0.6645	0.0521	0.4533	0.3190	0.4426
DataMaps-Easy	0.9972	0.0135	0.0771	0.0396	0.0928

Table B.1: Lexical and dialectal associations between toxicity in the original dataset (Davidson et al., 2017) and various filtered counterparts. Random, AFLite, and DataMaps all contain only 50% of the original data after filtering. (We could not perform downsampling on these datasets due to their heavily skewed label distribution.) Lower Pearson  $R$  correlation value indicates less superficial patterns in the dataset, thus are less biased. The easy subset gives the best results here are due to its severe imbalanced label distribution.

	Test		NOI		OI		ONI	
	Acc.↑	$F_1 \uparrow$	$F_1 \uparrow$	$\text{FPR}_{\text{NOI}} \downarrow$	$F_1 \uparrow$	$\text{FPR}_{\text{OI}} \downarrow$	$F_1 \uparrow$	$\text{FPR}_{\text{ONI}} \downarrow$
Original	96.37	97.81	96.42	25.00	99.86	57.14	99.57	63.64
LMIXIN-TOXTRIG	96.15	97.69	96.19	28.57	99.78	42.86	99.28	72.73

Table B.2: Lexical bias removal evaluation for debiasing methods. Original refers to the model trained over the full training set. The test set is further categorized into tweets that contained relevant TOXTRIG words.  $F_1$  indicates models’ performance while the false positive rate ( $\text{FPR}_*$ ) reflects models’ bias. The lower the  $\text{FPR}_*$  is, the less biased the model tend to be.

Debiasing Method	Test			
	$R_{AAE}$	Acc. $\uparrow$	$F_1$ $\uparrow$	$FPR_{AAE}$ $\downarrow$
Original	0.4079	96.37	97.81	24.76
LMIXIN-Dialect	-	96.48	97.88	22.86

Table B.3: Dialectal bias evaluation for all debiasing methods, on both in-distribution test set as well as out-of-distribution dialect and race priming test sets. In addition to accuracy and  $F_1$ , we report the false positive rate with respect to tweets in AAE ( $FPR_{AAE}$ ), reflecting dialectal bias (lower is less debiased). Each method is based on a RoBERTa-large classifier.

## Appendix C

### FEW-SHOT AAE-TO-WAE TRANSLATION

**Note that we do *not* recommend the following approach to build large scale parallel data for dialects, as discussed under ethical implications and limitations (§7).**

We use GPT-3 (Brown et al., 2020) to create a few-shot AAE-to-WAE translation system, using the following set of example translation pairs drawn from Spears (1998):

AAE: Get your triflin' ass out of here.

WAE: Get your trifling self out of here.

AAE: I saw his ass yesterday.

WAE: I saw him yesterday.

AAE: His ass is gonna get fried.

WAE: He is gonna get fried

AAE: Wassup, nigga?

WAE: What's up bro?

AAE: ⟨tweet⟩

WAE:

Note that Spears (Spears, 1998) refers to WAE as White language varieties, and deals with English prevalent in the United States.

We prepend the formatted example pairs to each AAE tweet in our training data, and generate the translation from GPT-3 using top-0.95 nucleus sampling with a temperature of 0.5. Prompts, formatting, and generation parameters were chosen based on manual inspection of the output.