

# Debiasing Word Embeddings with Nonlinear Geometry

Lu Cheng

Department of Computer Science,  
University of Illinois Chicago  
*lucheng@uic.edu*

Nayoung Kim and Huan Liu

School of Computing and Augmented Intelligence  
Arizona State University  
*{nkim48,huanliu}@asu.edu*

## Abstract

Debiasing word embeddings has been largely limited to individual and independent social categories. However, real-world corpora typically present multiple social categories that possibly correlate or intersect with each other. For instance, “hair weaves” is stereotypically associated with African American females, but neither African American nor females alone. Therefore, this work studies biases associated with multiple social categories: *joint biases* induced by the union of different categories and *intersectional biases* that do not overlap with the biases of the constituent categories. We first empirically observe that individual biases intersect *non-trivially* (i.e., over a one-dimensional subspace). Drawing from the intersectional theory in social science and the linguistic theory, we then construct an intersectional subspace to debias for multiple social categories using the nonlinear geometry of individual biases. Empirical evaluations corroborate the efficacy of our approach<sup>1</sup>.

## 1 Introduction

Due to the reliance on the large-scale text corpora for training, it has been observed that word embeddings are prone to express social biases inherent in the data (Bolukbasi et al., 2016; Caliskan et al., 2017). Prior research (e.g. Zhao et al., 2019; Bolukbasi et al., 2016) in debiasing word embeddings mitigates biases associated with individual social categories and treats each category in isolation. For example, the seminal Hard-Debiasing approach (Bolukbasi et al., 2016) identifies the bias direction of a category (e.g., gender) and then removes the direction from the target word such that it is equidistant to all groups (e.g., female and male at a binary level with gender) in the category. However, real-world training corpora typically present

<sup>1</sup>Data and implementation code can be downloaded at <https://github.com/GitHubLuCheng/Implementation-of-JoSEC-COLING-22>.




Figure 1: Illustrations of *joint* (i.e., union) and *intersectional* (i.e., intersect) biases using Gender and Race.

multiple social categories (e.g., gender and race), possibly with higher cardinality. These social categories can further correlate or intersect with each other (Thomas, 2004; Hancock, 2007). Despite the promising results, debiasing for individual social categories limits our understanding of the complex nature of social biases.

Alternatively, we might consider biases in the presence of multiple social categories (Foulds et al., 2020; Kearns et al., 2018; Cheng et al., 2022b,a). This can result in at least two scenarios, as depicted in Fig. 1. First, word embeddings can simultaneously present multiple biases that might *non-trivially correlate* with each other. For example, debiasing for gender can influence the results of racial bias. In the literature of social psychology, a number of works (e.g. Akrami et al., 2011; Bierly, 1985; Allport et al., 1954) studied the interrelationship between various biases, a.k.a. “generalized prejudice”. We refer to bias induced by the **union** of different social categories as **joint biases**. Second, a few recent works (e.g. Guo and Caliskan, 2020) detected **intersectional biases** in word embeddings, which is the bias that does not overlap with the biases of their constituent identities. For example, “hair weaves” is stereotypically associated with African American females (Ghavami and Peplau, 2013).

The primary goal of this work is to mitigate the two kinds of biases in word embeddings. There are several challenges: First, it is highly possible that different biases are nonlinearly correlated (Cheng

et al., 2022a). Simply taking a linear combination of individual bias subspaces (e.g., sum or mean of all bias subspaces) might lead to ineffective solutions that even amplify individual biases. Second, identifying a bias subspace typically needs predefined word sets related to target groups. These sets are curated by experts to most accurately represent each group and the associated social biases. This can be time-consuming and requires great human effort. Existing sets for identifying intersectional groups are extremely limited. They are small and exclusively used for the intersectionality of gender and race (Guo and Caliskan, 2020; Tan and Celis, 2019). Therefore, relying on predefined sets limits the use of debiasing approaches in practice.

To address these challenges, we first empirically observe that individual bias subspaces constructed via existing word sets intersect over a one-dimensional subspace. We then relate our findings to the intersectionality theory by Crenshaw (1989) and the linguistic theory introduced by Firth (1957). The result is a **hypothesis** that the bias subspace for multiple social categories roughly resides in the intersection of all individual bias subspaces. The proposed approach (coined as Joint and Intersectional Debiasing, JoSEC) departs from the linear correlation assumption and leverages the nonlinear geometry of subspace representations to learn an intersectional subspace. JoSEC does not need any additional human-coded defining sets for intersectional groups except for the defining sets for the constituent groups. We contribute to:

- a novel problem that considers biases associated with multiple social categories in mitigation;
- an effective approach for constructing the bias subspace without relying on the defining sets for intersectional categories; and
- demonstrations of the effectiveness of JoSEC via empirical evaluations on the benchmark datasets.

## 2 Related Work

Early efforts in debiasing word embeddings have been focused on gender. The seminal work by Bolukbasi et al. (2016) proposed a post-processing approach that projects gender-neutral words into a gender subspace identified by defining sets of gendered words such as *she*, *he*, *man*, *woman*. Gender bias can then be alleviated through hard-debiasing in which the bias components in non-

gendered words (e.g., *doctor*, *nurse*) are first removed and the gendered words are then centered and equalized. Manzini et al. (2019) further extended the hard-debiasing method to multi-class settings such as race. Given a corpus, one can also learn gender-neutral word embeddings by modifying the GloVe (Pennington et al., 2014) objective function (Zhao et al., 2018). This in-processing approach was further extended to a post-processing approach by Kaneko and Bollegala (2020), who suggested preserving gender-related information using autoencoder. Previous research as well as our work focuses on static word embeddings, however, stereotyped biases have also been found in contextualized word embeddings, e.g., (Zhao et al., 2019; Bordia and Bowman, 2019).

Despite the fruitful results, most prior works were found to remove biases superficially and fail to deliver gender-neutral embeddings (Gonen and Goldberg, 2019; Blodgett et al., 2020). Therefore, it is increasingly observable that existing bias removal techniques are insufficient to guarantee gender-neutral modeling. The majority of existing works were also criticized for not examining the impact of gender bias in real-world applications (Blodgett et al., 2020). In contrast to prior research focused on one form of bias in debiasing word embeddings, this work aims to provide a simple yet effective approach for bias mitigation in the presence of multiple bias forms. As human-like biases exist in the majority of word embeddings and debiasing approaches are unlikely to largely affect our results, we build our approach upon the seminal Hard-Debiasing algorithm by Bolukbasi et al. (2016). Future research is warranted to investigate other debiasing approaches.

For intersectional bias, most of the existing works are in social science and psychology literature, such as (Crenshaw, 1989; Kahn and Yoder, 1989; Hare-Mustin and Marecek, 1988). Comparatively fewer efforts can be found in the computer science field. One such work (Buolamwini and Gebu, 2018) examined the intersectional accuracy disparities in commercial gender classification systems. For contextualized word embeddings, May et al. (2019) and Tan and Celis (2019) measured the emergent intersectional biases of African American females using attributes presented in (Caliskan et al., 2017). In complement to WEAT, Guo and Caliskan (2020) proposed the Contextualized Embedding Association Test (CEAT) to measure the

intersectional bias. These methods detect intersectional bias using the defining and attribute sets related to intersectional groups.

In summary, this work complements prior research by mitigating biases related to multiple social categories. JoSEC does not rely on the human-coded word sets used to define the intersectional groups to identify the bias subspace. While this work focuses on non-contextualized embeddings, research (e.g., Guo and Caliskan, 2020; Lepori, 2020) has found contextualized word embeddings such as BERT (Devlin et al., 2019) display intersectional biases. We leave it for future exploration.

### 3 Preliminary

In this section, we briefly review the backbone model of this work: the seminal hard-debiasing method (Bolukbasi et al., 2016) and its extension. It consists of two steps: identifying the bias subspace and removing bias components.

#### 3.1 Identifying Bias Subspace

The individual bias subspace is identified by the *defining sets*, in which words represent different ends of the bias. For example, the defining sets of gender can be the gendered pronouns  $\{he, she\}$  and nouns  $\{man, woman\}$ . One can then identify the gender subspace  $\mathcal{B}$  by (1) subtracting the word embeddings of words in each defining set from the set’s mean, and (2) obtaining the  $K$  most significant components of the resulting vectors through a dimensionality-reduction method.

#### 3.2 Removing Bias Components

The next step is to apply the hard debiasing strategy to completely or partially remove the subspace components from the word embeddings. Hard-debiasing consists of two steps – *Neutralize* and *Equalize*. “Neutralize” removes bias components from non-gendered words (e.g., *doctor* and *nurse*); “Equalize” aims to center the gendered word embeddings (e.g., *she* and *he*) and equalize their bias components, such as the word pair  $\{man, woman\}$ .

Formally, given a bias subspace  $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K\}$ , where  $K$  denotes the number of principal components, we first compute the bias component  $\mathbf{w}_B$  of embedding  $\mathbf{w} \in \mathbb{R}^d$  in  $\mathcal{B}$  by

$$\mathbf{w}_B = \sum_{k=1}^K \langle \mathbf{w}, \mathbf{b}_k \rangle \mathbf{b}_k. \quad (1)$$

We then *neutralize* word embeddings by removing the bias component from non-gendered words:

$$\mathbf{w}' = \frac{\mathbf{w} - \mathbf{w}_B}{\|\mathbf{w} - \mathbf{w}_B\|}, \quad (2)$$

where  $\mathbf{w}'$  are the debiased word embeddings.

To “Equalize”, we debias the gendered words in a given equality set  $E$  by the following equation:

$$\mathbf{w}' = (\boldsymbol{\mu} - \boldsymbol{\mu}_B) + \sqrt{1 - \|\boldsymbol{\mu} - \boldsymbol{\mu}_B\|^2} \frac{\mathbf{w}_B - \boldsymbol{\mu}_B}{\|\mathbf{w}_B - \boldsymbol{\mu}_B\|}, \quad (3)$$

where  $\boldsymbol{\mu} = \frac{1}{|E|} \sum_{\mathbf{w} \in E} \mathbf{w}$  is the mean of embeddings of the words in the equality set  $E$ .  $\boldsymbol{\mu}_B$  denotes the bias component of  $\boldsymbol{\mu}$  in the identified bias subspace. It can be obtained via Eq. 1.

#### 3.3 Extending into Multi-Class Settings

In a multi-class setting (i.e., with more than two classes in a category, e.g., religion or race), the task inherently becomes non-linearly separable (Manzini et al., 2019). However, it is possible to linearly separate multiple classes based on the components of word embeddings. The multi-class bias subspace is then defined as follows: Given  $n$  defining sets of word embeddings  $\{D_1, D_2, \dots, D_n\}$ , the bias subspace  $\mathcal{B}$  is defined by the first  $K$  components of the following Principal Component Analysis (PCA) (Abdi and Williams, 2010) evaluation:

$$\mathcal{B} = \text{PCA} \left( \bigcup_{i=1}^n \bigcup_{\mathbf{w} \in D_i} \mathbf{w} - \boldsymbol{\mu}_i \right), \quad (4)$$

where  $\boldsymbol{\mu}_i = \frac{1}{|D_i|} \sum_{\mathbf{w} \in D_i} \mathbf{w}$  is the mean of word embeddings in set  $i$ .  $\bigcup$  denotes concatenation by rows. To remove multi-class bias, one can use the hard-debiasing method described in Sec. 3.2.

## 4 Method

Existing approaches for debiasing word embeddings work on individual categories, rendering incomplete measurement of various social biases (Hancock, 2007; Hurtado and Sinha, 2008). To account for biases associated with multiple social categories, we need to address the primary challenges of the potential non-linear correlations between biases and the difficulty of curating defining sets to identify the bias subspace for these categories. In this section, we introduce the proposed approach – JoSEC – for identifying such a subspace **without additional** human-coded defining set, which we refer to as the *intersectional subspace*.

## 4.1 Social Categories as “Cultural Contexts”

If individual biases are linearly correlated, we may construct the intersectional subspace by simply taking the sum or average of all individual bias subspaces. However, social biases are complex by nature as suggested by evidence in social science and psychology such as the “generalized prejudice”, i.e., generalized devaluing sentiments across different groups. To better quantify the potential non-linearity, we might first take a step back and revisit the development of single-word embeddings.

An influential position in the development of word embeddings holds that semantic representations for words can be derived through the patterns of lexical co-occurrence in language corpora. This is famously summarized by Firth (1957) as “you shall know a word by the company it keeps”. The central tenet is the idea that the sense of the target word could be inferred from its contexts, i.e., neighboring words within the sentence. Informed by this finding, we might assume that human-like biases follow a similar principle: a bias form w.r.t. some social category can be identified by its unique cultural contexts. In debiasing word embeddings, this indicates that the bias subspace (e.g., gender subspace) can be defined by the defining sets of words (e.g., *she*, *he*) that provide a specific “cultural context” for this bias form.

Naturally, when defining the intersectional subspace associated with multiple social categories, we might similarly consider each social category as a unique “cultural context” of the intersectional subspace. For example, the subspace of the intersectional group of gender and race is defined by both the gender subspace and the race subspace. That is, “Aisha” – a common name of an African American female – can be exclusively defined within a cultural context jointly determined by both gender and race whilst terms such as “hair weaves” should not depend on such context. Underpinning this assumption is the idea similar to the linguistic theory: each social category provides *unique context* to construct the intersectional subspace associated with multiple categories.

## 4.2 Geometry of Subspace Representation

Under the “social categories as cultural contexts” assumption, the intersectional subspace might have a fairly large intersection with each individual bias subspace. We empirically observe and hypothesize that the intersectional subspace of multiple

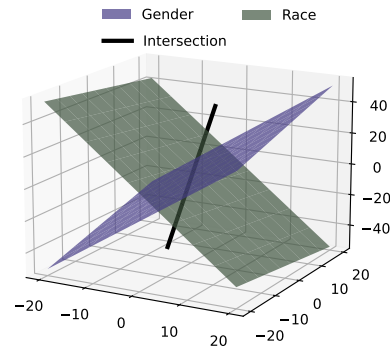


Figure 2: The geometry of Gender and Race bias subspaces and the intersectional subspace (Intersection).

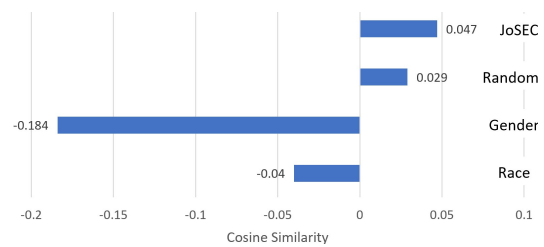


Figure 3: Cosine similarity between the ground-truth intersectional subspace and subspaces obtained via various approaches. “Random” denotes the result averaged over 10 similarity scores from 10 random vectors.

social categories should reside in all the subspaces representing the “cultural contexts” where the intersectional subspace is defined. Specifically, the intersectional subspace resides in the *intersection* of all individual bias subspaces and these subspaces should **intersect non-trivially**. This further implies that there exists a direction (one-dimension subspace) that is extremely close to all individual bias subspaces. We use this vector to represent the intersectional subspace. We propose the following hypothesis for identifying the bias subspace associated with multiple social categories:

**Hypothesis (Intersectional Hypothesis).** *The intersectional subspace  $\mathcal{B}_{sec}$  should reside in the intersection of  $\{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_N\}$ , where  $\mathcal{B}_i$  denotes the bias subspace of social category  $i$  and  $N$  is the number of considered social categories.*

Intersectional Hypothesis can be seen as operationalizing the intersectionality theory guided by Firth’s hypothesis.

### 4.2.1 Empirical Validation of the Intersectional Hypothesis

We empirically validate the intersectional hypothesis using the benchmark dataset L2-Reddit corpus (detailed in Sec. 5.1). The dataset includes the



defining sets for three social categories: race, gender, and religion, respectively. We first construct three individual bias subspaces for the target categories and each subspace has dimensions  $K \times d$ , where  $d$  denotes the dimension of word embeddings. For better visualization, we further project the  $d$ -dimensional subspace representations to 3D vectors using PCA.

We draw subspaces w.r.t. two randomly selected categories as 2-dimensional planes in Fig. 2. Result for the three categories is in Appendix A. We also visualize the corresponding intersectional subspaces identified by Eq. 7. We observe that the individual bias subspaces intersect roughly in a common direction with which the intersectional subspace approximately aligns. In addition, we use the defining sets for the intersectional groups of gender and race (e.g., African American Female) provided by WEAT (Caliskan et al., 2017) and Parada (2016) to construct the “ground-truth” intersectional bias subspace. We then calculate the cosine similarity between the ground-truth subspace and the (a) gender subspace, (b) race subspace, (c) random vector, and (d) intersectional subspace approximated by JoSEC, respectively. The similarity score ranges from -1 to 1, with -1 denoting the most dissimilar. Results in Fig. 3 suggest that JoSEC generates the intersectional subspace significantly more similar to the ground-truth subspace. Note that most of the similarity scores are close to 0. We believe this is in part because of the limitation of existing defining sets for intersectional identities, e.g., small in size. Both quantitative (Fig. 3) and qualitative (Fig. 2) analyses empirically justify the intersectional hypothesis.

#### 4.2.2 Identifying Intersectional Subspace

Under the *Intersectional Hypothesis*, we are essentially seeking the direction vector  $\mathcal{B}_{sec} = \hat{\mathbf{u}}$  that is “closest” to all individual bias subspaces. Let  $\mathbf{u}$  be a unit-length vector. We then reduce the problem of identifying intersectional subspace to the following optimization task:

$$\hat{\mathbf{u}} = \arg \min_{\|\mathbf{u}\|=1} \sum_{i=1}^N d(\mathbf{u}, \mathcal{B}_i)^2, \quad (5)$$

where  $d(\mathbf{u}, \mathcal{B}_i)$  is the shortest  $\ell_2$ -distance between the intersectional subspace and the bias subspace of social category  $i \in \{1, 2, \dots, N\}$ . Formally,

$$d(\mathbf{u}, \mathcal{B}_i) = \sqrt{\|\mathbf{u}\|^2 - \sum_{k=1}^K (\mathbf{u}^\top \mathbf{v}_{ik})^2}, \quad (6)$$

where  $\{\mathbf{v}_{i1}, \dots, \mathbf{v}_{iK}\}$  are the  $K$  principal components representing bias subspace  $\mathcal{B}_i$ . Eq. 5 can be reformulated as the following:

$$\hat{\mathbf{u}} = \arg \max_{\|\mathbf{u}\|=1} \sum_{i=1}^N \sum_{k=1}^K (\mathbf{u}^\top \mathbf{v}_{ik})^2. \quad (7)$$

Eq. 7 can be solved by taking the first principal component of  $\{\mathbf{v}_{ik}\}_{i=1, \dots, N; k=1, \dots, K}$ .

With the identified intersectional subspace  $\hat{\mathbf{u}}$ , we then follow Eq. 1-3 to remove the identified intersectional bias components from the target words (e.g., *hair weaves*).

## 5 Experiments

We validate the efficacy of the proposed intersectional subspace for debiasing word embeddings with multiple social categories. We answer the following research questions: How does JoSEC fare against baselines for mitigating (RQ. 1) *joint* biases and (RQ. 2) *intersectional* biases? How does JoSEC influence the (RQ. 3) utility of word embeddings in downstream tasks and (RQ. 4) biases in downstream tasks such as toxicity detection?

### 5.1 Language Corpus and Social Bias Data

Results for debiasing word embeddings are based on the commonly-used L2-Reddit corpus (Rabinovich et al., 2018), a collection of Reddit<sup>2</sup> posts and comments. It has been shown that the structural factor in user-generated content sites like Reddit make them less welcoming to marginalized populations (Bender et al., 2021). The initial biased word embeddings are obtained by training word2vec on approximately 56 million sentences. It includes three social categories: *Gender*, *Race*, and *Religion*. We use vocabularies from (Bolukbasi et al., 2016) and (Caliskan et al., 2017) as the defining and attribute sets for gender. Word sets for race and religion are the same lexicons used in (Manzini et al., 2019).

Following (Guo and Caliskan, 2020), we consider the intersectionality of race and gender for the evaluation of intersectional debiasing due to the limited data availability. In particular, there are in total  $3 \times 2$  (3 racial classes and 2 gender classes) intersectional social groups: {*African American male*, *African American female*, *European American male*, *European American female*,

<sup>2</sup><https://www.reddit.com/>

*Mexican American male, Mexican American female*}. The defining sets for these groups are provided by WEAT (Caliskan et al., 2017) and Parada (2016), including frequent given names that represent group membership. The intersectional attribute sets identified through human workers’ validation are provided by Ghavami and Peplau (2013). For example, one common given name included in the defining set of African American females is “Aisha” and some related stereotyped bias include attributes “aggressive” and “dark-skinned”. For the complete list of given names and attributes, please refer to Appendix B. Note that in contrast to common debiasing approaches, JoSEC does not need these defining sets for joint and intersectional debiasing. We use them for evaluation only.

## 5.2 Experimental Setup

We briefly summarize the experimental settings, including the baselines, downstream tasks, and evaluation metrics.

### 5.2.1 Baselines

We are not aware of any existing approaches for joint or intersectional debiasing for word embeddings. Therefore, we adapt the hard-debiasing approach such that it works in both tasks. We also consider two debiasing strategies (i.e., SUM and MEAN) that impose a linear assumption to construct the bias subspace for multiple social categories. All compared approaches only differ in the subspace construction, the bias mitigation follows the same procedure in the hard-debiasing approach. Baselines are detailed below.

- *Hard\_Seq*. The hard-debiasing method extended to the *joint* debiasing task. In particular, *Hard\_Seq* sequentially debiases for individual social categories. As the order might influence the results of *Hard\_Seq*, we experiment with all potential sequences and report the best results.
- *Hard\_Insec*. The hard-debiasing method extended to the *intersectional* debiasing task. In particular, *Hard\_Insec* uses the human-coded defining sets for the intersectionality groups of gender and race in (Guo and Caliskan, 2020) to construct the intersectional subspace.
- *SUM*. Its bias subspace is constructed by summing up the subspaces of individual biases, i.e.,  $\mathcal{B}_{SUM} = \sum_{i=1}^N \mathcal{B}_i$ .
- *MEAN*. Its bias subspace is constructed by averaging over the subspaces of individual biases, i.e.,  $\mathcal{B}_{MEAN} = \frac{1}{N} \sum_{i=1}^N \mathcal{B}_i$ .

### 5.2.2 Downstream Tasks

**Utility (RQ. 3):** To examine the influence of JoSEC on the utility of word embeddings, we perform several standard downstream tasks following (Manzini et al., 2019). They are the CoNLL 2003 shared tasks (Sang and De Meulder, 2003), including NER tagging, POS (part-of-speech) tagging, and POS chunking. There are two evaluation paradigms: replacing the biased embeddings with the debiased ones or retraining the model on debiased embeddings. We only report results for one setting and the other can be found in Appendix C.

**Extrinsic Bias (RQ. 4):** While we work on mitigating biases in the pre-trained resource, i.e., *intrinsic bias*, recent research (e.g., Seraphina et al., 2021; Delobelle et al., 2021) presents interesting findings about the biases in downstream tasks enabled by word embeddings, i.e., *extrinsic bias*. Therefore, we further investigate how the debiased word embeddings influence biases in a common downstream NLP task, toxicity detection. Particularly, we consider the Kaggle Challenge of Jigsaw Unintended Bias in Toxicity Classification<sup>3</sup> and examine biases against gender, race, and religion. The Perspective API’s Jigsaw dataset has both toxicity and identity annotations. The training and test splits are the same as the original data. Please refer to Appendix D for detailed experimental settings for toxicity classification.

### 5.2.3 Evaluation Metrics

For debiasing tasks, we use the mean average cosine similarity (MAC) to quantify the intrinsic bias, as suggested in (Manzini et al., 2019). Given a set of target word embeddings  $\mathcal{S}$  of words with a specific form of social bias (e.g., *Jew, Christian, Muslim*) and a set of attribute sets  $\mathcal{A} = \{A_1, A_2, \dots, A_N\}$ ,  $A_j$  consists of embeddings of words  $\mathbf{a}$  (e.g., *violent, terrorist, uneducated*) that should not be associated with any word in  $\mathcal{S}$ . Let  $f(\cdot)$  be a function that computes the mean cosine distance between  $S_i \in \mathcal{S}$  and  $\mathbf{a} \in A_j$ :

$$f(S_i, A_j) = \frac{1}{|A_j|} \sum_{\mathbf{a} \in A_j} \cos(S_i, \mathbf{a}), \quad (8)$$

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

Table 1: MACs ( $\uparrow$ ) of all approaches for joint debiasing. “Re”, “Ra”, and “Ge” denotes “Religion”, “Race”, and “Gender”, respectively. “Re→Ra→Ge” denotes the order for sequential debiasing. We experiment with all potential sequences and report the best results. A larger value is more desired.

Model \ MAC	Gender	Race	Religion	Total
Biased	0.623	0.892	0.859	2.374
Hard_Seq (Re→Ra→Ge)	0.656	0.888	<b>0.937</b>	2.481
Hard_Seq (Ra→Re→Ge)	0.654	<b>0.929</b>	0.868	2.451
SUM	0.598	0.870	0.900	2.368
MEAN	0.657	0.872	0.862	2.391
JoSEC	<b>0.703</b>	0.914	0.917	<b>2.534</b>

where  $\cos(S_i, \mathbf{a}) = 1 - \frac{S_i \cdot \mathbf{a}}{\|S_i\|_2 \cdot \|\mathbf{a}\|_2}$ . MAC is then computed by

$$\text{MAC}(\mathcal{S}, \mathcal{A}) = \frac{1}{|\mathcal{S}| |\mathcal{A}|} \sum_{S_i \in \mathcal{S}} \sum_{A_j \in \mathcal{A}} f(S_i, A_j). \quad (9)$$

A larger MAC score denotes a greater bias removal.

For the downstream task that examines the utility, we report F1 scores from using biased word embeddings as well as the changes of F1 ( $\Delta$  F1), Precision ( $\Delta$  Precision), and Recall ( $\Delta$  Recall) after using debiased word embeddings. To check the statistical significance, we also perform a paired  $t$ -test on the distribution of average cosine distance used to compute MAC and student  $t$ -test for the results of downstream tasks. For the downstream task that examines the extrinsic bias, we use two common evaluation metrics: False Positive Equality Difference (FPED) and False Negative Equality Difference (FNED). We report Total (FPED+FNED) scores from using biased word embeddings and  $\Delta$  Total,  $\Delta$  FPED, and  $\Delta$  FNED after using debiased word embeddings. Unless otherwise noted, all the results below are statistically significant at level 0.05.

### 5.3 Results

We first present results (averaged over 5 repetitions) for **RQ. 1-4** and then discuss our findings.

#### 5.3.1 RQ. 1: Joint Debiasing

The joint debiasing task seeks to simultaneously mitigate biases induced by the union of all social categories. We report MACs w.r.t. each individual bias as well as the total bias, which is computed over all considered social categories. All the best results are highlighted.

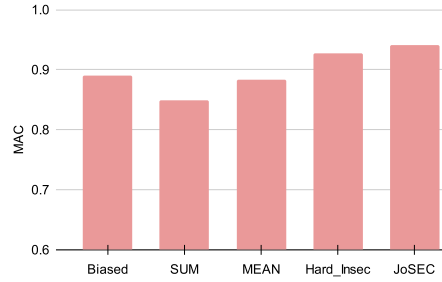


Figure 4: MACs ( $\uparrow$ ) of all compared approaches for the intersectional debiasing task. A larger value is more desired.

From the results in Table 1, we can observe that (1) the best results w.r.t. debiasing for individual categories are achieved by various approaches and the proposed approach (*JoSEC*) outperforms all baselines regarding reducing the total amount of bias (Total). This suggests that it is challenging to debias for all social categories simultaneously and the proposed intersectional subspace is effective for joint debiasing. (2) The linear solutions to subspace construction (i.e., *SUM* and *MEAN*) are not as effective as sequential debiasing. This empirically validates our hypothesis that social biases are non-linearly correlated and multiple social categories should intersect non-trivially.

(3) Of particular interest is that when debiasing sequentially, the bias mitigation performance w.r.t. the first category appears to be the most effective. For example, *Hard\_Seq* (Re→Ra→Ge) shows the best MAC of “Religion”. Further, by comparing the results in the second row (*Biased*) with those in the third (*Hard\_Seq*), the racial MAC of sequential debiasing decreases (i.e., more biases) whilst applying hard-debiasing to racial bias alone can actually lead to higher MACs. These findings might suggest that different biases are *interacting* with each other, i.e., they are correlated. Future research is warranted to examine the bias correlations and their influence on the debiasing approaches.

#### 5.3.2 RQ. 2: Intersectional Debiasing

To evaluate the intersectional bias removal performance of all the compared methods, we use the human-coded defining and attribute sets associated with the intersectionality of gender and race. Note that only *Hard\_Insec* used the human-coded defining sets to construct the bias subspace. All other approaches (i.e., *SUM*, *MEAN*, and *JoSEC*) use subspaces of individual biases to construct the subspace for multiple social categories. MACs of all

Table 2: Utility of biased and debiased word embeddings in NER Tagging (NER), POS Tagging (POS-T), and POS Chunking (POS-C) tasks, under the *Embedding Matrix Replacement* paradigm. Word embeddings are debiased by hard-debiasing with different subspaces.  $\Delta$  denotes the change before and after debiasing.

Tasks	Hard_Seq			Hard_Insec			SUM			MEAN			JoSEC		
	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C
Biased F1	0.96	0.99	1.00	0.96	0.99	1.00	0.96	0.99	1.00	0.96	0.99	1.00	0.96	0.99	1.00
$\Delta$ F1	-0.01	0.01	0.00	-0.01	0.00	0.00	-0.01	0.01	0.00	-0.01	0.00	0.00	-0.02	0.00	0.00
$\Delta$ Precision	-0.01	0.00	0.00	-0.02	0.00	0.00	-0.02	0.00	0.00	-0.02	0.00	0.00	-0.03	0.00	0.00
$\Delta$ Recall	-0.01	0.03	0.02	-0.02	0.01	0.02	-0.02	0.02	0.01	-0.02	0.01	0.01	-0.03	0.01	0.01

Table 3: Gender, racial, and religious biases in toxicity classification using biased and debiased word embeddings. Total=FPED+FNED ( $\downarrow$ ). A smaller bias score denotes less bias. A negative  $\Delta$  indicates reduced bias.

Bias	Hard_Seq			Hard_Insec			SUM			MEAN			JoSEC		
	Gender	Race	Religion	Gender	Race	Religion	Gender	Race	Religion	Gender	Race	Religion	Gender	Race	Religion
Biased Total	1.27	0.64	0.37	1.27	0.64	0.37	1.27	0.64	0.37	1.27	0.64	0.37	1.27	0.64	0.37
$\Delta$ Total	-0.04	0.04	0.05	-0.05	-0.01	0.01	0.00	0.04	-0.01	0.00	-0.01	0.01	-0.05	0.02	0.02
$\Delta$ FPED	-0.09	0.03	0.00	-0.05	0.03	-0.02	-0.03	0.03	-0.01	0.06	0.04	0.01	-0.04	0.03	0.00
$\Delta$ FNED	0.05	0.01	0.05	0.00	-0.04	0.02	0.03	0.01	0.01	-0.06	-0.04	0.00	0.00	-0.01	0.02

methods are presented in Fig. 4.

We make the following observations: (1) *JoSEC* is most effective for mitigating the intersectional biases w.r.t. gender and race. The fact that it outperforms *Hard\_Insec* manifests the potential to leverage the nonlinear geometry of subspace representation to construct the subspace for intersectional bias. This is encouraging as it suggests that we may not need defining sets for intersectional groups, which are often inaccessible and challenging to collect; (2) *SUM* and *MEAN* are outperformed by *Biased*, indicating that simply taking a linear combination of all the individual bias subspaces can aggravate the intersectional biases in word embeddings. The violation of the linear correlation assumption has a negative influence on the performance of hard-debiasing.

### 5.3.3 RQ. 3: Downstream Utility

This research question aims to investigate the effects of various debiasing strategies on the semantic utility of word embeddings in standard NLP tasks. We consider NER Tagging, POS Tagging, and POS Chunking, following (Manzini et al., 2019).

Results for embedding matrix replacement are in Table 2. We can observe that debiasing word embeddings with multiple social categories only slightly changes the semantic utility. We also perform the student *t* test and further testify that these differences are statistically insignificant. These results imply that debiasing for multiple categories using the hard-debiasing method does not have a significant influence on the semantic utility of word embeddings. This applies to both joint and intersectional debiasing.

### 5.3.4 RQ. 4: Bias in Toxicity Classification

This experiment further examines the extrinsic biases of word embeddings, particularly, in toxicity classification. Results are shown in Table 3. We observe that the debiased word embeddings have little influence on reducing biases in the downstream task, indicating no reliable correlation between the *intrinsic bias* and *extrinsic bias*. This aligns well with findings shown in (Seraphina et al., 2021). Aware of the importance of mitigating biases in downstream applications, it is critical to extend this work to joint and intersectional debiasing focused on extrinsic measures of biases in the future.

## 6 Discussions

This work studies **joint biases** induced by the union of multiple categories and the **intersectional biases** that do not overlap with biases of the constituent categories. Challenges arise from the potential nonlinearity between different biases and the difficulty of curating human-coded word lists for identifying intersectional bias subspace. We first empirically showed that different biases intersect *non-trivially*. Informed by the intersectionality theory and the linguistic theory by Firth, we propose a simple yet effective approach (JoSEC) for constructing the intersectional bias subspace using the nonlinear geometry of bias subspaces. JoSEC can reduce intrinsic bias without losing the semantic utility of word embeddings. The broad result of this research is that it is critical to consider biases associated with multiple social categories given the complex nature of human-like biases.

We do note several limitations of this study. First, the empirical observation of the nonlinearity be-



tween different biases needs more rigorous theoretical proof/evidence. A potential result would be an in-depth analysis of bias correlations in a variety of NLP tasks. Questions such as “how are biases correlated? Negative or positive correlation?” might be investigated. Second, research should be conducted to study how to reduce joint and intersectional biases in contextualized word embeddings such as BERT and GPT-3. Third, JoSEC focuses on intrinsic measures, which remain good descriptive metrics for computational social science (Seraphina et al., 2021). However, it might not be relied on to mitigate biases in downstream applications. As real-world scenarios can be more challenging and complicated (e.g., biases induced in the deployment), it is necessary to extend this research to extrinsic measures of biases, which might vary in different applications. Meanwhile, research needs to focus on a more rigorous and transparent data collection process (e.g., recording demographic and identity information of annotators) (Cheng et al., 2021) to help reduce downstream biases. Challenging sets to measure application bias such as (Röttger et al., 2021) need to be created to test the robustness of debiasing methods.

### Ethic Statement

We heavily rely on existing bias and fairness metrics, which certainly have no guarantee of *unbiased* word embeddings or machine learning models. In fact, most metrics can only be considered an indicator of bias at most (Seraphina et al., 2021), especially since significant limitations w.r.t. these metrics have been found (Garrido-Muñoz et al., 2021). Therefore, we urge practitioners not to rely on these debiased word embeddings alone, but also at least consider bias mitigation in specific downstream tasks. Further, we did not discuss many other negative impacts of language models that practitioners should consider, such as high energy consumption or not including all stakeholders in the design phase (Bender et al., 2021).

### Acknowledgements

This work is supported by the Office of Naval Research (ONR) under the grant N00014-21-1-4002 and National Science Foundation (NSF) under the grants #2036127 and #2114789. Lu Cheng is also supported by the startup fund provided by the Department of Computer Science at the University of Illinois Chicago. The views, opinions, and/or

findings expressed are the authors’ and should not be interpreted as representing the official views or policies of the funding agency.

### References

- Hervé Abdi and Lynne J Williams. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459.
- Nazar Akrami, Bo Ekehammar, and Robin Bergh. 2011. Generalized prejudice: Common and specific components. *Psychological Science*, 22(1):57–59.
- Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. 1954. The nature of prejudice.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Margaret M Bierly. 1985. Prejudice toward contemporary outgroups as a generalized attitude. *Journal of Applied Social Psychology*, 15(2):189–199.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. In *ACL*, pages 5454–5476.
- 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 *NeurIPS*.
- Shikha Bordia and Samuel R Bowman. 2019. Identifying and reducing gender bias in word-level language models. In *NAACL SRW*.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FACCT*, pages 77–91. PMLR.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Lu Cheng, Suyu Ge, and Huan Liu. 2022a. Toward understanding bias correlations for mitigation in nlp. *arXiv preprint arXiv:2205.12391*.
- Lu Cheng, Ahmadreza Mosallanezhad, Yasin N Silva, Deborah L Hall, and Huan Liu. 2022b. Bias mitigation for toxicity detection via sequential decisions. In *SIGIR*, pages 1750–1760.
- Lu Cheng, Kush R Varshney, and Huan Liu. 2021. Socially responsible ai algorithms: Issues, purposes, and challenges. *Journal of Artificial Intelligence Research*, 71:1137–1181.

- Kimberlé Crenshaw. 1989. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *u. Chi. Legal f.*, page 139.
- Pieter Delobelle, Ewoenam Kwaku Tokpo, Toon Calders, and Bettina Berendt. 2021. Measuring fairness with biased rulers: A survey on quantifying biases in pretrained language models. *arXiv preprint arXiv:2112.07447*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *AIES*, pages 67–73.
- J Firth. 1957. A synopsis of linguistic theory 1930–1955 in studies in linguistic analysis, philological society.
- James R Foulds, Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. 2020. An intersectional definition of fairness. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 1918–1921. IEEE.
- Ismael Garrido-Muñoz, Arturo Montejó-Ráez, Fernando Martínez-Santiago, and L Alfonso Ureña-López. 2021. A survey on bias in deep nlp. *Applied Sciences*, 11(7):3184.
- Negin Ghavami and Letitia Anne Peplau. 2013. An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly*, 37(1):113–127.
- Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In *NAACL-HLT*, pages 609–614.
- Wei Guo and Aylin Caliskan. 2020. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *AIES*.
- Ange-Marie Hancock. 2007. When multiplication doesn’t equal quick addition: Examining intersectionality as a research paradigm. *Perspectives on politics*, 5(1):63–79.
- Rachel T Hare-Mustin and Jeanne Marecek. 1988. The meaning of difference: Gender theory, postmodernism, and psychology. *American psychologist*, 43(6):455.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Aída Hurtado and Mrinal Sinha. 2008. More than men: Latino feminist masculinities and intersectionality. *Sex roles*, 59(5-6):337–349.
- Arnold S Kahn and Janice D Yoder. 1989. The psychology of women and conservatism: Rediscovering social change. *Psychology of women quarterly*, 13(4):417–432.
- Masahiro Kaneko and Danushka Bollegala. 2020. Autoencoding improves pre-trained word embeddings. In *COLING*.
- Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, pages 2564–2572. PMLR.
- Michael A Lepori. 2020. Unequal representations: Analyzing intersectional biases in word embeddings using representational similarity analysis. In *COLING*.
- Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. 2019. Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in word embeddings. In *NAACL-HLT*, pages 615–621.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. In *NAACL-HLT*.
- Maryann Parada. 2016. Ethnolinguistic and gender aspects of latino naming in chicago: Exploring regional variation. *Names*, 64(1):19–35.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Ella Rabinovich, Yulia Tsvetkov, and Shuly Wintner. 2018. Native language cognate effects on second language lexical choice. *Transactions of the Association for Computational Linguistics*, 6:329–342.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. Hatecheck: Functional tests for hate speech detection models. In *ACL*, pages 41–58.
- Erik F Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*.
- Goldfarb-Tarrant Seraphina, Marchant Rebecca, Sanchez Ricardo Muñoz, Pandya Mugdha, and Lopez Adam. 2021. Intrinsic bias metrics do not correlate with application bias. In *ACL*.
- Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. In *NeurIPS*.

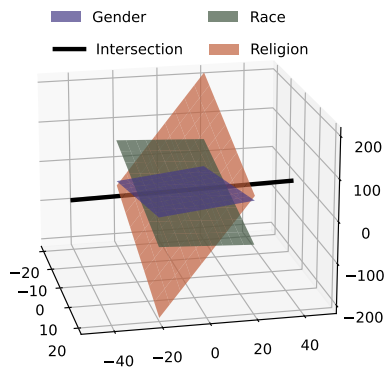


Figure 5: The geometry of Gender, Race, and Religion bias subspaces and the intersectional subspaces (Intersection).

Veronica G Thomas. 2004. The psychology of black women: Studying women’s lives in context. *Journal of Black Psychology*, 30(3):286–306.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender bias in contextualized word embeddings. In *NAACL*, pages 629–634.

Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning gender-neutral word embeddings. In *EMNLP*.

## A Additional Empirical Results for Validating the Intersectional Hypothesis

Following the same procedure in Sec. 4.2.1, we draw subspaces w.r.t. all of three categories as 2-dimensional planes in Figure 5. We observe similar results to those w.r.t. two categories.

## B Stimuli for Identifying Intersectional Groups and Biases

The stimuli used to construct the “ground-truth” intersectional bias subspace are from (Caliskan et al., 2017) and (Parada, 2016). We follow (Guo and Caliskan, 2020) to identify the names of intersectional groups and corresponding attributes associated with the intersectional biases. This process is further verified by human subjects (Ghavami and Peplau, 2013). We use this validation set to evaluate the quality of the intersectional subspace constructed by JoSEC. In particular, the ground-truth intersectional subspace is used in the empirical validation of the Intersectional Hypothesis (Appendix A) and comparisons of the debiasing performance between JoSEC and Hard\_Insec.

Names to identify intersectional groups:

- **African American females:** Aisha, Keisha, Lakisha, Latisha, Latoya, Malika, Nichelle, Shereen, Tamika, Tanisha, Yolanda, Yvette
- **African American males:** Alonzo, Alphonse, Hakim, Jamal, Jamel, Jerome, Leroy, Lionel, Marcellus, Terrence, Tyrone, Wardell
- **European American females:** Carrie, Colleen, Ellen, Emily, Heather, Katie, Megan, Melanie, Nancy, Rachel, Sarah, Stephanie
- **European American males:** Andrew, Brad, Frank, Geoffrey, Jack, Jonathan, Josh, Matthew, Neil, Peter, Roger, Stephen
- **Mexican American females:** Adriana, Alejandra, Alma, Brenda, Carolina, Iliana, Karina, Liset, Maria, Mayra, Sonia, Yesenia
- **Mexican American males:** Alberto, Alejandro, Alfredo, Antonio, César, Jesús, José, Juan, Miguel, Pedro, Rigoberto, Rogelio

Attributes indicating intersectional biases:

- **Intersectional biases of African American females:** aggressive, athletic, bigbutt, confident, darkskinned, friedchicken, ghetto, loud, overweight, promiscuous, unfeminine, unintelligent, unrefined
- **Intersectional biases of African American males:** athletic, criminals, dangerous, darkskinned, gangsters, hypersexual, lazy, loud, poor, rapper, tall, unintelligent, violent
- **Intersectional Biases of European American females:** arrogant, attractive, blond, ditsy, emotional, feminine, highstatus, intelligent, materialistic, petite, racist, rich, submissive, tall
- **Intersectional biases of European American males:** allAmerican, arrogant, attractive, blond, high-status, intelligent, leader, privileged, racist, rich, sexist, successful, tall
- **Intersectional biases of Mexican American females:** cook, curvy, darkskinned, feisty, hardworker, loud, maids, promiscuous, sexy, short, uneducated, unintelligent
- **Intersectional biases of Mexican American males:** aggressive, arrogant, darkskinned,

day-laborer, drunks, hardworker, illegal-immigrant, jealous, macho, poor, promiscuous, short, uneducated, unintelligent, violent

### C Additional Experimental Results for RQ. 3

This section presents results for the semantic utility of debiased word embeddings under the *Model Retraining* paradigm. In particular, we show Precision, Recall, F1 scores w.r.t. NER Tagging, POS Tagging, and POS Chunking, respectively. We retrain the model with the debiased word embeddings. As shown in Table 4, we observe similar results to those using Embedding Replacement. Together with the results in Sec. 5.3.3, we show that debiasing for multiple social categories using the hard-debiasing method does not have a significant influence on the semantic utility of word embeddings. This applies to both joint and intersectional debiasing.

### D Experimental Setup for Toxicity Classification in RQ. 4

We employ LSTM (Long Short-Term Memory) (Hochreiter and Schmidhuber, 1997) as the toxicity classifier, which takes the input of a variety of word embeddings considered in this work. We then measure the biases using two commonly adopted metrics (Dixon et al., 2018): False Positive Equality Difference (FPED) and False Negative Equality Difference (FNED). FNED/FPED is defined as the sum of deviations of group-specific False Negative Rates (FNRs)/False Positive Rates (FPRs) from the overall FNR/FPR. Given  $N$  demographic groups (e.g., female and male in gender) and we denote each group as  $G_{i \in \{1, \dots, N\}}$ , FNED and FPED are calculated as:

$$\begin{aligned} FNED &= \sum_{i \in \{1, \dots, N\}} |FNR - FNR_{G_i}|, \\ FPED &= \sum_{i \in \{1, \dots, N\}} |FPR - FPR_{G_i}|. \end{aligned} \quad (10)$$

where  $FNR_{G_i}$  denotes the FNR calculated over group  $G_i$  and FNR is calculated over the entire training set. A debiased model is expected to have similar FNR and FPR for different groups belonging to the same identity, therefore, smaller FNED and FPED are desired. Ideally, the sum of FNED and FPED is close to zero.



Table 4: Utility of biased and debiased word embeddings in NER Tagging (NER), POS Tagging (POS-T), and POS Chunking (POS-C) tasks, under the *Model Retraining* paradigm. Word embeddings are debiased by hard-debiasing with different subspaces.  $\Delta$  denotes the change before and after debiasing.

Tasks	Hard_Seq			Hard_Insec			SUM			MEAN			JoSEC		
	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C	NER	POS-T	POS-C
Biased F1	0.97	1.00	0.99	0.97	1.00	0.99	0.97	1.00	0.99	0.97	1.00	0.99	0.97	1.00	0.99
$\Delta$ F1	-0.01	0.01	0.00	-0.01	0.00	0.00	-0.01	0.01	0.00	-0.00	0.00	0.00	-0.00	0.00	0.00
$\Delta$ Precision	-0.00	0.00	0.00	-0.01	0.00	0.00	-0.01	0.00	0.00	0.00	0.00	0.00	-0.00	0.00	0.00
$\Delta$ Recall	-0.02	0.03	0.02	-0.02	0.01	0.02	-0.01	0.02	0.01	-0.01	0.01	0.01	-0.00	0.01	0.01