

Conceptor Debiasing of Word Representations Evaluated on WEAT

Saket Karve

Lyle Ungar

João Sedoc

Department of Computer & Information Science

University of Pennsylvania

Philadelphia, PA 19104

{saketk, ungar, joao} @cis.upenn.edu

Abstract

Bias in word embeddings such as Word2Vec has been widely investigated, and many efforts made to remove such bias. We show how to use *conceptors debiasing* to post-process both traditional and contextualized word embeddings. Our conceptor debiasing can simultaneously remove racial and gender biases and, unlike standard debiasing methods, can make effect use of heterogeneous lists of biased words. We show that conceptor debiasing diminishes racial and gender bias of word representations as measured using the Word Embedding Association Test (WEAT) of Caliskan et al. (2017).

1 Introduction

Word embeddings capture distributional similarities and thus inherit demographic stereotypes (Bolukbasi et al., 2016). Such embedding biases tend to track statistical regularities such as the percentage of people with a given occupation (Nikhil Garg and Zou, 2018) but sometimes deviate from them (Bhatia, 2017). Recent work has shown that gender bias exists in contextualized embeddings (Wang et al., 2019; May et al., 2019).

Here, we provide a quantitative analysis of bias in traditional and contextual word embeddings and introduce a method of mitigating bias (i.e., debiasing) using *the debiasing conceptor*, a clean mathematical representation of subspaces that can be operated on and composed by logic-based manipulations (Jaeger, 2014). Specifically, conceptor negation is a soft damping of the principal components of the target subspace (e.g., the subset of words being debiased) (Liu et al., 2019b) (See Figure 1.)

Key to our method is how it treats word-association lists (sometimes called target lists), which define the bias subspace. These lists include pre-chosen words associated with a target



(a) The original space



(b) After applying the debiasing conceptor

Figure 1: BERT word representations of the union of the set of contextualized word representations of *relatives*, *executive*, *wedding*, *salary* projected on to the first two principal components of the WEAT gender first names, which capture the primary component of gender. Note how the debiasing conceptor collapses *relatives* and *wedding*, and *executive* and *salary* once the bias is removed.

demographic group (often referred to as a “protected class”). For example, *he / she* or *Mary / John* have been used for gender (Bolukbasi et al., 2016). More generally, conceptors can combine multiple subspaces defined by word lists. Unlike most current methods, conceptor debiasing uses a

soft, rather than a hard projection.

We test the debiasing conceceptor on a range of traditional and contextualized word embeddings¹ and examine whether they remove stereotypical demographic biases. All tests have been performed on English word embeddings.

This paper contributes the following:

- Introduces *debiasing conceptors* along with a formal definition and mathematical relation to the Word Embedding Association Test.
- Demonstrates the effectiveness of the debiasing conceceptor on both traditional and contextualized word embeddings.

2 Related Work

NLP has begun tackling the problems that inhibit the achievement of fair and ethical AI (Hovy and Spruit, 2016; Friedler et al., 2016), in part by developing techniques for mitigating demographic biases in models. In brief, a *demographic bias* is a difference in model output based on gender (either of the data author or of the content itself) or selected demographic dimension (“protected class”) such as race. Demographic biases manifest in many ways, ranging from disparities in tagging and classification accuracy depending on author age and gender (Hovy, 2015; Dixon et al., 2018), to over-amplification of demographic differences in language generation (Yatskar et al., 2016; Zhao et al., 2017), to diverging implicit associations between words or concepts within embeddings or language models (Bolukbasi et al., 2016; Rudinger et al., 2018).

Here, we are concerned with the societal bias towards protected classes that manifests in prejudice and stereotypes (Bhatia, 2017). Greenwald and Banaji (1995); implicit attitudes such that “introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feeling, thought, or action toward social objects.” Bias is often quantified in people using the Implicit Association Test (IAT) (Greenwald et al., 1998). The IAT records subjects response times when asked to pair two concepts. Smaller response times occur in concepts subjects perceive to be similar versus pairs of concepts they perceive to be different. A well known example is where subjects were asked to associate black

¹Previous work has shown that debiasing methods can have different effects on different word embeddings (Kiritchenko and Mohammad, 2018).

and white names with “pleasant” and “unpleasant” words. A significant racial bias has been found in many populations. Later, Caliskan et al. (2017) formalized the Word Embedding Association Test (WEAT), which replaces reaction time with word similarity to give a bias measure that does not require use of human subjects. May et al. (2019) extended WEAT to the Sentence Embedding Association Test (SEAT); however, in this paper we instead use token-averaged representations over a corpus.

Debiasing Embeddings. The simplest way to remove bias is to project out a bias direction. For example, Bolukbasi et al. (2016) identify a “gender subspace” using lists of gendered words and then remove the first principal component of this subspace. Wang et al. (2019) used both data augmentation and debiasing of Bolukbasi et al. (2016) to mitigate bias found in ELMo and showed improved performance on coreference resolution. Our work is complementary, as debiasing conceptors can be used in place of hard-debiasing.

Bolukbasi et al. (2016) also examine a soft debiasing method, but find that it does not perform well. In contrast, our debiasing conceceptor does a successful soft damping of the relevant principal components. To understand why, we first introduce the conceceptor method for capturing the “bias subspaces”, next formalize bias, and then show WEAT in matrix notation.

2.1 Conceptors

As in Bolukbasi et al. (2016), our aim is to identify the “bias subspace” using a set of target words, \mathcal{Z} and Z is their corresponding word embeddings. A conceceptor matrix, C , is a regularized identity map (in our case, from the original word embeddings to their biased versions) that minimizes

$$\|Z - CZ\|_F^2 + \alpha^{-2} \|C\|_F^2. \quad (1)$$

where α^{-2} is a scalar parameter.²

To describe matrix conceptors, we draw heavily on (Jaeger, 2014; He and Jaeger, 2018; Liu et al., 2019b,a). C has a closed form solution:

$$C = \frac{1}{k} ZZ^\top \left(\frac{1}{k} ZZ^\top + \alpha^{-2} I \right)^{-1}. \quad (2)$$

Intuitively, C is a soft projection matrix on the linear subspace where the word embeddings Z have

²Note that the conceceptor and WEAT literature disagree on notation and we follow WEAT. In conceceptor notation, the matrix Z would be denoted as X .

the highest variance. Once C has been learned, it can be ‘negated’ by subtracting it from the identity matrix and then applied to any word embeddings to shrink the bias directions.

Conceptors can represent laws of Boolean logic, such as NOT \neg , AND \wedge and OR \vee . For two conceptors C and B , we define the following operations:

$$\neg C := \mathbf{I} - C, \quad (3)$$

$$C \wedge B := (C^{-1} + B^{-1} - \mathbf{I})^{-1} \quad (4)$$

$$C \vee B := \neg(\neg C \wedge \neg B) \quad (5)$$

Among these Boolean operations, two are critical for this paper: the NOT operator for debiasing, and the OR operation \vee for multi-list (or multi-category) debiasing. It can be shown that if C and B are of equal sizes, then $C \vee B$ is the conceptor computed from the union of the two sets of sample points from which C and B are computed (Jaeger, 2014); this is not true if they are of different sizes.

Negated Conceptor. Given that the conceptor, C , represents the subspace of maximum bias, we want to apply the negated conceptor, NOT C (see Equation 3) to an embedding space and remove its bias. We call NOT C the *debiasing conceptor*. More generally, if we have K conceptors, C_i derived from K different word lists, we call NOT $(C_1 \vee \dots \vee C_K)$ a debiasing conceptor. The negated conceptor matrix has been used in the past on a complete vocabulary to increase the semantic richness of its word embeddings; Liu et al. (2018) showed that the negated conceptor gave better performance on semantic similarity and downstream tasks than the hard debiasing method of Mu and Viswanath (2018).

As shown in Liu et al. (2018), the negated conceptor approach does a soft debiasing by shrinking each principal component of the covariance matrix of the target word embeddings ZZ^\top . The shrinkage is a function of the conceptor hyper-parameter α and the singular values σ_i of ZZ^\top : $\frac{\alpha^{-2}}{\sigma_i + \alpha^{-2}}$.

3 Formalizing Bias

We follow the formal definition of Lu et al. (2018), where given a class of word sets \mathcal{D} and a scoring function s , the bias of s under the concept(s) tested by \mathcal{D} , written $\mathcal{B}_s(\mathcal{D})$, is the expected difference in scores assigned to expected absolute bias across class members,

$$\mathcal{B}_s(\mathcal{D}) \triangleq \mathbb{E}_{D \in \mathcal{D}} |\mathcal{B}_s(D)|.$$

This naturally gives rise to a large set of concepts and scoring functions.

3.1 Word Embedding Association Test

The Word Embeddings Association Test (WEAT), as proposed by Caliskan et al. (2017), is a statistical test analogous to the Implicit Association Test (IAT) (Greenwald et al., 1998) which helps quantify human biases in textual data. WEAT uses the cosine similarity between word embeddings, which is analogous to the reaction time when subjects are asked to pair two concepts they find similar in the IAT. WEAT considers two sets of target words and two sets of attribute words of equal size. The null hypothesis is that there is no difference between the two sets of target words and the sets of attribute words in terms of their relative similarities measured as the cosine similarity between the embeddings. For example, consider the target sets as words representing *Career* and *Family* and let the two sets of attribute words be *Male* and *Female*, in that order. The null hypothesis states that *Career* and *Family* are equally similar (mathematically, in terms of the mean cosine similarity between the word representations) to each of the words in the *Male* and *Female* word lists.

The WEAT test statistic measures the differential association of the two sets of target words with the attribute. The ‘‘effect size’’ is a normalized measure of how separated the two distributions are.

To ground this, we cast WEAT in our formulation where \mathcal{X} and \mathcal{Y} are two sets of target words, (concretely, \mathcal{X} might be *Career* words and \mathcal{Y} *Family* words) and \mathcal{A} , \mathcal{B} are two sets of attribute words (\mathcal{A} might be *female* names and \mathcal{B} *male* names) assumed to associate with the bias concept(s). WEAT is then ³

$$\begin{aligned} s(\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}) &= \frac{1}{|\mathcal{X}|} \left[\sum_{x \in \mathcal{X}} \left[\sum_{a \in \mathcal{A}} s(x, a) - \sum_{b \in \mathcal{B}} s(x, b) \right] \right. \\ &\quad \left. - \sum_{y \in \mathcal{Y}} \left[\sum_{a \in \mathcal{A}} s(y, a) - \sum_{b \in \mathcal{B}} s(y, b) \right] \right], \end{aligned}$$

where $s(x, y) = \cos(\text{vec}(x), \text{vec}(y))$ and $\text{vec}(x) \in \mathbb{R}^k$ is the k -dimensional word embedding for word x . Note that for this definition of

³We assume that there is no overlap between any of the sets \mathcal{X} , \mathcal{Y} , \mathcal{A} , and \mathcal{B} .

WEAT, the cardinality of the sets must be equal, so $|\mathcal{A}| = |\mathcal{B}|$ and $|\mathcal{X}| = |\mathcal{Y}|$. Our conceptor formulation given below relaxes this assumption.

To motivate our conceptor formulation, we further generalize WEAT to capture the covariance between the target word and the attribute word embeddings. First, let X, Y, A and B be matrices whose columns are word embeddings corresponding to the words in the sets $\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}$, respectively (i.e. the two sets of target words and two sets of attribute words, respectively). To formally define this, without loss of generality choose \mathcal{X} , let $X = [x_i]_{i \in I}$ where for i in an index set I with cardinality $|\mathcal{X}|$ and $x_i = \text{vec}(x)$ where the word x is indexed at the i th value of the index set.⁴ We can then write WEAT as,

$$\begin{aligned} & \|X^T A - X^T B - (Y^T A - Y^T B)\|_F \\ &= \|(X - Y)^T (A - B)\|_F, \end{aligned}$$

where $\|\cdot\|_F$ is the Frobenius norm. If the embeddings are unit length, then GWEAT is the same as $|\mathcal{X}|$ times WEAT.⁵

Suppose we want to mitigate bias by applying the $k \times k$ bias mitigating matrix, $G = -C$, which optimally removes bias from any matrix of word embeddings. We select G to minimize

$$\begin{aligned} & \|(G(X - Y))^T G(A - B)\|_F, \\ &= \|(X - Y)^T G^T G(A - B)\|_F. \end{aligned}$$

Since the conceptor, C , is calculated using the word embeddings of $\mathcal{Z} = \mathcal{X} \cup \mathcal{Y}$, the negated conceptor will mitigate the variance from the target sets, which hopefully identifies the most important bias directions.

4 Embeddings

For context-independent embeddings, we used off-the-shelf Fasttext subword embeddings⁶, which were trained with subword information on the Common Crawl (600B tokens), the GloVe embeddings⁷ trained on Wikipedia and Gigaword and word2vec⁸ trained on roughly 100 billion

⁴To clarify, in our notation $x_i \in \mathbb{R}^k$ and $x \in \mathcal{X}$.

⁵Our generalization of WEAT is different from Swinger et al. (2018).

⁶<https://dl.fbaipublicfiles.com/fasttext/vectors-english/crawl-300d-2M-subword.zip>.

⁷<https://nlp.stanford.edu/projects/glove/>

⁸

words from a Google News dataset. The embeddings used are not centered and normalized to unit length as in Bolukbasi et al. (2016).

For contextualized embeddings, we used ELMo small which was trained on the 1 Billion Word Benchmark, approximately 800M tokens of news crawl data from WMT 2011.⁹ We also experimented with the state-of-the-art contextual model ‘‘BERT-Large, Uncased’’ which has 24-layer, 1024-hidden, 16-heads, 340M parameters. BERT is trained on the BooksCorpus (0.8B words) and Wikipedia (2.5B words). We used the last four hidden layers of BERT. We used the Brown Corpus for the word contexts to create instances of the ELMo and BERT embeddings. Embeddings of English words only have been used for all the tests.

5 WEAT Debiasing Experiments

As described in section 3.1, WEAT assumes as its null hypothesis that there is no relative bias between the pair of concepts defined as the target words and attribute words. In our experiments, we measure the effect size (the WEAT score normalized by the standard deviation of differences of attribute words w.r.t target words) (d) and the one-sided p-value of the permutation test. A higher absolute value of effect size indicates larger bias between words in the target set with respect to the words in the attribute set. We would like the absolute value of the effect size to be zero. Since the p-value measures the likelihood that a random permutation of the attribute words would produce at least the observed test statistic, it should be high (at least 0.05) to indicate lack of bias in the positive direction.

Conceptually, the conceptor should be a soft projection matrix on the linear subspace representing the bias direction. For instance, the subspace representing gender must consist of words which are specific to or in some sense related to gender.

A gender word list might be a set of pronouns which are specific to a particular gender such as *he / she* or *himself / herself* and gender specific words representing relationships like *brother / sister* or *uncle / aunt*. We test conceptor debiasing both using the list of such pronouns used by

⁹https://s3-us-west-2.amazonaws.com/allennlp/models/elmo/2x1024_128_2048cnn_1xhighway/elmo_2x1024_128_2048cnn_1xhighway_weights.hdf5

Embedding	Subspace	Without Debiasing		Mu et al.		Bolukbasi et al.		Conceptor Negation	
		d	p	d	p	d	p	d	p
Glove	Pronouns	1.78	0.00	1.81	0.00	1.24	0.01	0.13	0.40
	Extended List			1.86	0.00	1.24	0.01	0.36	0.26
	Propernouns			1.74	0.00	1.24	0.01	0.78	0.07
	All			1.75	0.00	1.20	0.01	0.35	0.27
	OR			NA	NA	NA	NA	-0.51	0.81
word2vec	Pronouns	1.81	0.00	1.79	0.00	1.55	0.00	1.09	0.02
	Extended List			1.79	0.00	1.59	0.00	1.38	0.00
	Propernouns			1.70	0.0	1.59	0.0	1.45	0.00
	All			1.71	0.00	1.56	0.00	1.40	0.00
	OR			NA	NA	NA	NA	0.84	0.05
Fasttext	Pronouns	1.67	0.00	1.70	0.0	1.45	0.00	0.95	0.04
	Extended List			1.70	0.0	1.47	0.00	0.84	0.04
	Propernouns			0.86	0.06	1.47	0.00	0.85	0.06
	All			0.82	0.05	1.14	0.01	0.81	0.06
	OR			NA	NA	NA	NA	0.24	0.33

Table 1: Gender Debiasing non-contextualized embeddings: (Career, Family) vs (Male, Female)

Embedding	Subspace	Without Debiasing		Mu et. al.		Conceptor Negation	
		d	p	d	p	d	p
ELMo	Pronouns	1.79	0.0	1.79	0.00	0.70	0.10
	Extended List			1.79	0.00	0.06	0.46
	Propernouns			1.79	0.00	-0.61	0.89
	All			1.79	0.00	-0.28	0.73
	OR			NA	NA	-0.85	0.96
BERT	Pronouns	1.21	0.01	1.21	0.01	1.31	0.00
	Extended List			1.27	0.00	1.33	0.01
	Propernouns			1.27	0.01	0.92	0.04
	All			1.27	0.01	0.63	0.13
	OR			NA	NA	0.97	0.02

Table 2: Gender Debiasing Contextualized embeddings: (Career, Family) vs (Male, Female)

Embedding	Subspace	Without Debiasing		Mu et al.		Bolukbasi et al		Conceptor Negation	
		d	p	d	p	d	p	d	p
Glove	Pronouns	1.09	0.02	0.89	0.04	-0.53	0.85	1.04	0.01
	Extended List			1.07	0.02	-0.60	0.86	-0.52	0.83
	Propernouns			1.04	0.02	-0.56	0.86	0.20	0.33
	All			1.03	0.02	-0.53	0.82	0.18	0.35
	OR			NA	NA	NA	NA	-0.48	0.82
Word2vec	Pronouns	1.00	0.02	0.89	0.03	-1.09	0.99	1.10	0.01
	Extended List			1.00	0.03	-1.14	1.00	-0.49	0.82
	Propernouns			0.88	0.04	-1.17	1.00	0.33	0.27
	All			0.90	0.04	-1.07	0.99	0.25	0.34
	OR			NA	NA	NA	NA	-0.47	0.81
Fasttext	Pronouns	1.19	0.01	1.08	0.01	0.18	0.35	-0.36	0.76
	Extended List			0.71	0.08	0.21	0.353	0.73	0.09
	Propernouns			0.12	0.43	0.15	0.40	-0.47	0.80
	All			0.038	0.47	0.20	0.32	-0.50	0.84
	OR			NA	NA	NA	NA	-0.46	0.78

Table 3: Gender Debiasing non-contextualized embeddings: (Math, Arts) vs (Male, Female)

Embedding	Subspace	Without Debiasing		Mu et. al.		Conceptor Negation	
		d	p	d	p	d	p
ELMo	Pronouns	0.94	0.02	0.94	0.03	-0.03	0.38
	Extended List			0.95	0.02	0.27	0.29
	Propernouns			0.94	0.02	0.85	0.05
	All			0.94	0.04	0.87	0.05
	OR			NA	NA	0.53	0.13
BERT	Pronouns	0.23	0.777	0.23	0.79	0.15	0.15
	Extended List			0.16	0.82	0.06	0.53
	Propernouns			0.16	0.82	0.75	0.08
	All			0.16	0.85	0.43	0.24
	OR			NA	NA	-0.07	0.59

Table 4: Gender Debiasing contextualized embeddings: (Math, Arts) vs (Male, Female)

Embedding	Subspace	Without Debiasing		Mu et al.		Bolukbasi et al.		Conceptor Negation	
		d	p	d	p	d	p	d	p
Glove	Pronouns	1.34	0.0	1.23	0.01	-0.46	0.819	-0.20	0.66
	Extended List			1.27	0.00	-0.51	0.83	0.93	0.04
	Propernouns			1.21	0.011	-0.48	0.839	0.65	0.10
	All			1.21	0.00	-0.45	0.81	0.68	0.10
	OR			NA	NA	NA	NA	0.60	0.12
Word2vec	Pronouns	1.16	0.01	1.09	0.02	-0.46	0.80	0.45	0.21
	Extended List			1.20	0.01	-0.50	0.80	0.59	0.13
	Propernouns			1.08	0.02	-0.55	0.86	0.69	0.10
	All			1.08	0.02	-0.46	0.80	0.66	0.13
	OR			NA	NA	NA	NA	0.09	0.45
Fasttext	Pronouns	1.48	0.00	1.51	0.00	0.88	0.04	0.93	0.03
	Extended List			0.85	0.04	0.85	0.04	1.36	0.00
	Propernouns			1.01	0.03	0.85	0.05	0.75	0.08
	All			0.98	0.03	0.88	0.03	0.89	0.05
	OR			NA	NA	NA	NA	0.89	0.05

Table 5: Gender Debiasing non-cotextualized embeddings: (Science, Arts) vs (Male, Female)

Embedding	Subspace	Without Debiasing		Mu et. al.		Conceptor Negation	
		d	p	d	p	d	p
ELMo	Pronouns	1.32	0.0	1.31	0.00	0.41	0.22
	Extended List			1.32	0.005	0.52	0.24
	Propernouns			1.38	0.00	1.28	0.00
	All			1.34	0.00	0.92	0.03
	OR			NA	NA	0.82	0.05
BERT	Pronouns	-0.91	0.88	-0.91	0.87	-1.23	0.97
	Extended List			-0.90	0.91	-1.10	0.99
	Propernouns			-0.90	0.92	-0.93	0.92
	All			-0.90	0.90	-0.38	0.70
	OR			NA	NA	0.97	0.02

Table 6: Gender Debiasing cotextualized embeddings: (Science, Arts) vs (Male, Female)

Caliskan et al. (2017) and using a more comprehensive list of gender-specific words that includes gender specific terms related to occupations, relationships and other commonly used words such as

prince / princess and *host / hostess*¹⁰. We further tested conceptor debiasing using male and female

¹⁰<https://github.com/uclanlp/corefBias>,
https://github.com/uclanlp/gn_glove

names such as *Aaron / Alice* or *Chris / Clary*¹¹. We also tested our method with the combination of all lists. The combination of the subspace was done in two ways - either by taking the union of all word lists or by applying the OR operator on the three conceptor matrices computed independently.

The subspace for racial bias was determined using list of European American and African American names.

We tested target pairs of Science vs. Arts, Math vs. Arts, and Career vs. Family word lists with the attribute of the male vs. female names to test gender debiasing. Similarly, we examined European American names vs. African American names as target pairs with the attribute of pleasant vs. unpleasant to test racial debiasing.

Our findings indicate that expanded lists give better debiasing for word embeddings; however, the results are not as clear for contextualized embeddings. The OR operator on two conceptors describing subspaces of pronouns/nouns and names generally outperforms a union of these words. This further motivates the use of the debiasing conceptor.

5.1 Racial Debiasing Results

Embedding	Original		Conceptor Negation	
	d	p	d	p
GloVe	1.35	0.00	0.69	0.01
word2vec	-0.27	0.27	-0.55	0.72
Fasttext	0.41	0.04	-0.27	0.57
ELMo	1.37	0.00	-0.45	0.20
BERT	0.92	0.00	0.36	0.61

Table 7: Racial Debiasing: (European American Names, African American Names) vs (Pleasant, Unpleasant). d is the effect size, which we want to be close to 0 and p is the p-value, which we want to be larger than 0.05.

Table 7 summarizes the effect size (d) and the one-sided p-value we obtained by running WEAT on each of the word embeddings for racial debiasing. In this experiment we used the same setup as Caliskan et al. (2017) and compare attribute Words of European American / African American names with target words “pleasant” and “unpleasant”. In Table 7 we see that racial bias is mitigated

¹¹<https://www.cs.cmu.edu/Groups/AI/areas/nlp/corpora/names/>

in all cases aside from GloVe. Furthermore, for word2vec the associational bias is not significant. We also found that the conceptor nearly always outperforms the hard debiasing methods of Mu and Viswanath (2018) and Bolukbasi et al. (2016).

5.2 Gender Debiasing Results

Tables 1, 3 and 5 show the results obtained on gender debiasing between attribute words of “Family” and “Career”, “Math” and “Arts” and “Science” and “Arts” with the target words “Male” and “Female” respectively for the traditional word embeddings. We show the results for all the word representations; however, the method of Bolukbasi et al. (2016) can only be applied to standard word embeddings.¹² We show the results when embeddings are debiased using conceptors computed using different subspaces. It can be seen in the tables that the bias for the conceptor negated embeddings is significantly less than that of the original embeddings. In the tables, the conceptor debiasing method is compared with the hard-debiasing technique proposed by Mu and Viswanath (2018) where the first principal component of the subspace from the embeddings is completely project off. The debiasing conceptor outperforms the hard debiasing technique in almost all cases. Note that the OR operator can not be used with the hard debiasing technique and thus is not reported.

Similarly, Tables 2, 4 and 6 show a comparison of the effect size and p-value using the hard debiasing technique and conceptor debiasing on conceptualized embeddings. It can be seen that conceptor debiasing generally outperforms other methods in mitigating (has a small absolute value) bias with the ELMo embeddings for all the subspaces. The results are less clear for BERT as observed in Table 6, which we will discuss in the following section. Note that combining all subspaces gives a significant reduction in the effect size.

5.3 Discussion of BERT Results

One of our most surprising findings is that unlike ELMo, the bias in BERT according to WEAT is less consistent than other word representations; WEAT effect sizes in BERT vary largely across different layers. Furthermore, the debiasing conceptor occasionally creates reverse bias in BERT, suggesting that tuning of the hyper-parameter α

¹²The concurrent work of Wang et al. (2019) was not available in time for us to compare with this method.

may be required. Another possibility is that BERT is capturing multiple concepts, and the presumption that the target lists are adequately capturing gender or racial attributes is incorrect. This suggests that further study into word lists is called for, along with visualization and end-task evaluation. It should also be noted that our results are in line with those from [May et al. \(2019\)](#).

6 Retaining Semantic Similarity

In order to understand if the debiasing conceceptor was harming the semantic content of the word embeddings, we examined conceceptor debiased embedding for semantic similarity tasks. As done in [Liu et al. \(2018\)](#) we used the seven standard word similarity test set and report Pearson’s correlation. The word similarity sets are: the RG65 ([Rubenstein and Goodenough, 1965](#)), the WordSim-353 (WS) ([Finkelstein et al., 2002](#)), the rare-words (RW) ([Luong et al., 2013](#)), the MEN dataset ([Bruni et al., 2014](#)), the MTurk ([Radinsky et al., 2011](#)), the SimLex-999 (SimLex) ([Hill et al., 2015](#)), and the SimVerb-3500 ([Gerz et al., 2016](#)). Table 8 shows that conceceptors help in preserving and at times increasing the semantic information in the embeddings. It should be noted that these tasks can not be applied to contextualized embeddings such as ELMo and BERT. So, we do not report these results.

	GloVe		word2vec		Fasttext	
	Orig.	CN	Orig.	CN	Orig.	CN
RG65	76.03	70.92	74.94	78.58	85.87	85.94
WS	73.79	75.17	69.34	69.34	78.82	77.44
RW	51.01	55.25	55.78	56.04	62.17	62.48
MEN	80.13	80.10	77.07	77.85	83.64	82.64
MTurk	69.16	71.17	68.31	67.68	72.45	71.34
SimLex	40.76	45.85	44.27	46.05	50.55	50.78
SimVerb	28.42	34.51	36.54	37.33	42.75	42.72

Table 8: Word Similarity comparison with conceceptor debiased embeddings using all gender words as conceceptor subspace.

7 Conclusion

We have shown that the debiasing conceceptor can successfully debias word embeddings, outperforming previous state-of-the art ‘hard’ debiasing methods. Best results are obtained when lists are broken up into subsets of ‘similar’ words (pronouns, professions, names, etc.), and separate conceceptors are learned for each subset and then OR’d. Conceceptors for different protected subclasses such

as gender and race can be similarly OR’d to jointly debias.

Contextual embeddings such as ELMo and BERT, which give a different vector for each word token, work particularly well with conceceptors, since they produce a large number of embeddings; however, further research on tuning conceceptors for BERT needs to be done. Finally, we note that embedding debiasing may leave bias which is undetected by measures such as WEAT [Gonen and Goldberg \(2019\)](#); thus, all debiasing methods should be tested on end-tasks such as emotion classification and co-reference resolution.

References

- Sudeep Bhatia. 2017. The semantic representation of prejudice and stereotypes. *Cognition*, 164:46–60.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357.
- E. Bruni, N. K. Tran, and M. Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49(1):1–47.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 67–73. ACM.
- L. Finkelstein, E. Gabrilovich, Y. Matias, E. Rivlin, Z. Solan, G. Wolfman, and E. Ruppin. 2002. Placing search in context: the concept revisited. *ACM Transactions on Information Systems*, 20(1):116–131.
- Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236*.
- D. Gerz, I. Vulic, F. Hill, R. Reichart, and A. Korhonen. 2016. SimVerb-3500: a large-scale evaluation set of verb similarity. In *Proceedings of the EMNLP 2016*, pages 2173–2182.
- 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 *North American Chapter of the Association for Computational Linguistics (NAACL)*.

- Anthony G Greenwald and Mahzarin R Banaji. 1995. Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological review*, 102(1):4.
- Anthony G Greenwald, Debbie E McGhee, and Jordan LK Schwartz. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464.
- X. He and H. Jaeger. 2018. [Overcoming catastrophic interference using conceptor-aided back-propagation](#). In *International Conference on Learning Representations*.
- F. Hill, R. Reichart, and A. Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Dirk Hovy. 2015. Demographic factors improve classification performance. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 752–762.
- Dirk Hovy and Shannon L Spruit. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 591–598.
- H. Jaeger. 2014. [Controlling recurrent neural networks by conceptors](#). Technical report, Jacobs University Bremen.
- Svetlana Kiritchenko and Saif Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53.
- T. Liu, J. Sedoc, and L. Ungar. 2018. Correcting the common discourse bias in linear representation of sentences using conceptors. In *Proceedings of ACM-BCB- 2018 Workshop on BioCreative/OHNL Challenge, Washington, D.C., 2018*.
- T. Liu, L. Ungar, and J. Sedoc. 2019a. Continual learning for sentence representations using conceptors. In *Proceedings of the NAACL HLT 2019*.
- T. Liu, L. Ungar, and J. Sedoc. 2019b. [Unsupervised post-processing of word vectors via conceptor negation](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-2019), Honolulu*.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2018. Gender bias in neural natural language processing. *arXiv preprint arXiv:1807.11714*.
- M. Luong, R. Socher, and C. D. Manning. 2013. Better word representations with recursive neural networks for morphology. In *Proceedings of the CoNLL 2013*.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- J. Mu and P. Viswanath. 2018. All-but-the-top: Simple and effective postprocessing for word representations. In *International Conference on Learning Representations*.
- Dan Jurafsky Nikhil Garg, Londa Schiebinger and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *PNAS*.
- K Radinsky, E. Agichtein, E. Gabrilovich, and S. Markovitch. 2011. A word at a time: Computing word relatedness using temporal semantic analysis. In *Proceedings of the 20th International World Wide Web Conference*, pages 337–346, Hyderabad, India.
- H. Rubenstein and J. B. Goodenough. 1965. Contextual correlates of synonymy. *Communications of the ACM*, 8(10):627–633.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. *arXiv preprint arXiv:1804.09301*.
- Nathaniel Swinger, Maria De-Arteaga, Neil Thomas Heffernan IV, Mark D. M. Leiserson, and Adam Tauman Kalai. 2018. [What are the biases in my word embedding?](#) *CoRR*, abs/1812.08769.
- Tianlu Wang, Jieyu Zhao, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender bias in contextualized word embeddings. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5534–5542.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989.