

It’s All in the Name: Mitigating Gender Bias with Name-Based Counterfactual Data Substitution

Rowan Hall Maudslay¹ Hila Gonen² Ryan Cotterell¹ Simone Teufel¹

¹ Department of Computer Science and Technology, University of Cambridge

² Department of Computer Science, Bar-Ilan University

{rh635, rdc42, sht25}@cam.ac.uk hilagnn@gmail.com

Abstract

This paper treats gender bias latent in word embeddings. Previous mitigation attempts rely on the operationalisation of gender bias as a projection over a linear subspace. An alternative approach is *Counterfactual Data Augmentation* (CDA), in which a corpus is duplicated and augmented to remove bias, e.g. by swapping all inherently-gendered words in the copy. We perform an empirical comparison of these approaches on the English Gigaword and Wikipedia, and find that whilst both successfully reduce direct bias and perform well in tasks which quantify embedding quality, CDA variants outperform projection-based methods at the task of drawing non-biased gender analogies by an average of 19% across both corpora. We propose two improvements to CDA: *Counterfactual Data Substitution* (CDS), a variant of CDA in which potentially biased text is randomly substituted to avoid duplication, and the Names Intervention, a novel name-pairing technique that vastly increases the number of words being treated. CDA/S with the Names Intervention is the only approach which is able to mitigate indirect bias: following debiasing, previously biased words are significantly less clustered according to gender (cluster purity is reduced by 49%), thus improving on the state-of-the-art for bias mitigation.

1 Introduction

Gender bias describes an inherent prejudice against a gender, captured both by individuals and larger social systems. Word embeddings, a popular machine-learned semantic space, have been shown to retain gender bias present in corpora used to train them (Caliskan et al., 2017). This results in gender-stereotypical vector analogies à la Mikolov et al. (2013), such as *man:computer programmer :: woman:homemaker* (Bolukbasi et al., 2016), and such bias has been shown to materialise in a variety of downstream tasks, e.g. coreference resolution (Rudinger et al., 2018; Zhao et al., 2018).

By operationalising gender bias in word embeddings as a linear subspace, Bolukbasi et al. (2016) are able to debias with simple techniques from linear algebra. Their method successfully mitigates **direct bias**: *man* is no longer more similar to *computer programmer* in vector space than *woman*. However, the structure of gender bias in vector space remains largely intact, and the new vectors still evince **indirect bias**: associations which result from gender bias between not explicitly gendered words, for example a possible association between *football* and *business* resulting from their mutual association with explicitly masculine words (Gonen and Goldberg, 2019). In this paper we continue the work of Gonen and Goldberg, and show that another paradigm for gender bias mitigation proposed by Lu et al. (2018), Counterfactual Data Augmentation (CDA), is also unable to mitigate indirect bias. We also show, using a new test we describe (non-biased gender analogies), that WED might be removing too much gender information, casting further doubt on its operationalisation of gender bias as a linear subspace.

To improve CDA we make two proposals. The first, Counterfactual Data Substitution (CDS), is designed to avoid text duplication in favour of substitution. The second, the Names Intervention, is a method which can be applied to either CDA or CDS, and treats bias inherent in first names. It does so using a novel name pairing strategy that accounts for both name frequency and gender-specificity. Using our improvements, the clusters of the most biased words exhibit a reduction of cluster purity by an average of 49% across both corpora following treatment, thereby offering a partial solution to the problem of indirect bias as formalised by Gonen and Goldberg (2019). Additionally, although one could expect that the debiased embeddings might suffer performance losses in computational linguistic tasks, our embeddings remain useful for at least two such tasks, word similarity and sentiment classification (Le and Mikolov, 2014).

2 Related Work

The measurement and mitigation of gender bias relies on the chosen operationalisation of gender bias. As a direct consequence, how researchers choose to operationalise bias determines both the techniques at one’s disposal to mitigate the bias, as well as the yardstick by which success is determined.

2.1 Word Embedding Debiasing

One popular method for the mitigation of gender bias, introduced by Bolukbasi et al. (2016), measures the genderedness of words by the extent to which they point in a **gender direction**. Suppose we embed our words into \mathbb{R}^d . The fundamental assumption is that there exists a linear subspace $B \subset \mathbb{R}^d$ that contains (most of) the gender bias in the space of word embeddings. (Note that B is a direction when it is a single vector.) We term this assumption the **gender subspace hypothesis**. Thus, by basic linear algebra, we may decompose any word vector $\mathbf{v} \in \mathbb{R}^d$ as the sum of the projections onto the bias subspace and its complement: $\mathbf{v} = \mathbf{v}_B + \mathbf{v}_{\perp B}$. The (implicit) operationalisation of gender bias under this hypothesis is, then, the magnitude of the bias vector $\|\mathbf{v}_B\|_2$.

To capture B , Bolukbasi et al. (2016) first construct N sets D_i , each of which contains a pair of words that differ in their gender but that are otherwise semantically equivalent (using a predefined set of gender-definitional pairs). For example, $\{man, woman\}$ would be one set and $\{husband, wife\}$ would be another. They then compute the average empirical covariance matrix

$$C = \sum_{i=1}^N \frac{1}{|D_i|} \sum_{w \in D_i} (\vec{w} - \mu_i)(\vec{w} - \mu_i)^\top \quad (1)$$

where μ_i is the mean embedding of the words in D_i , then B is taken to be the space spanned by the top k eigenvectors of C associated with the largest eigenvalues. Bolukbasi et al. set $k = 1$, and thus define a gender direction.

Using this operationalisation of gender bias, Bolukbasi et al. go on to provide a linear-algebraic method (Word Embedding Debiasing, WED, originally “hard debiasing”) to remove gender bias in two phases: first, for non-gendered words, the gender direction is removed (“neutralised”). Second, pairs of gendered words such as *mother* and *father* are made equidistant to all non-gendered words (“equalised”). Crucially, under the gender subspace

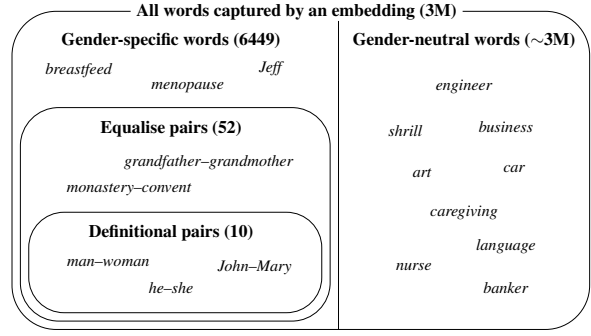


Figure 1: Word sets used by WED with examples

hypothesis, it is only necessary to identify the subspace B as it is possible to perfectly remove the bias under this operationalisation using tools from numerical linear algebra.

The method uses three sets of words or word pairs: 10 definitional pairs (used to define the gender direction), 218 gender-specific seed words (expanded to a larger set using a linear classifier, the complement of which is neutralised in the first step), and 52 equalise pairs (equalised in the second step). The relationships among these sets are illustrated in Figure 1; for instance, gender-neutral words are defined as all words in an embedding that are not gender-specific.

Bolukbasi et al. find that this method results in a 68% reduction of stereotypical analogies as identified by human judges. However, bias is removed only insofar as the operationalisation allows. In a comprehensive analysis, Gonen and Goldberg (2019) show that the original structure of bias in the WED embedding space remains intact.

2.2 Counterfactual Data Augmentation

As an alternative to WED, Lu et al. (2018) propose Counterfactual Data Augmentation (CDA), in which a text transformation designed to invert bias is performed on a text corpus, the result of which is then appended to the original, to form a new bias-mitigated corpus used for training embeddings. Several interventions are proposed: in the simplest, occurrences of words in 124 gendered word pairs are swapped. For example, ‘the woman cleaned the kitchen’ would (counterfactually) become ‘the man cleaned the kitchen’ as *man-woman* is on the list. Both versions would then together be used in embedding training, in effect neutralising the *man-woman* bias.

The *grammar intervention*, Lu et al.’s improved intervention, uses coreference information to veto

swapping gender words when they corefer to a proper noun.¹ This avoids *Elizabeth ... she ... queen* being changed to, for instance, *Elizabeth ... he ... king*. It also uses POS information to avoid ungrammaticality related to the ambiguity of *her* between personal pronoun and possessive determiner. In the context, ‘her teacher was proud of her’, this results in the correct sentence ‘his teacher was proud of *him*’.

3 Improvements to CDA

We prefer the philosophy of CDA over WED as it makes fewer assumptions about the operationalisation of the bias it is meant to mitigate.

3.1 Counterfactual Data Substitution

The duplication of text which lies at the heart of CDA will produce debiased corpora with peculiar statistical properties unlike those of naturally occurring text. Almost all observed word frequencies will be even, with a notable jump from 2 directly to 0, and a type–token ratio far lower than predicted by Heaps’ Law for text of this length. The precise effect this will have on the resulting embedding space is hard to predict, but we assume that it is preferable not to violate the fundamental assumptions of the algorithms used to create embeddings.

As such, we propose to apply substitutions probabilistically (with 0.5 probability), which results in a non-duplicated counterfactual training corpus, a method we call **Counterfactual Data Substitution (CDS)**. Substitutions are performed on a per-document basis in order to maintain grammaticality and discourse coherence. This simple change should have advantages in terms of naturalness of text and processing efficiency, as well as theoretical foundation.

3.2 The Names Intervention

Our main technical contribution in this paper is to provide a method for better counterfactual augmentation, which is based on bipartite-graph matching of names. Instead of Lu et. al.’s (2018) solution of not treating words which corefer to proper nouns in order to maintain grammaticality, we propose an explicit treatment of first names. This is because we note that as a result of not swapping the gender of words which corefer with proper nouns, CDA

¹We interpret Lu et al.’s (2018) phrase “cluster” to mean “coreference chain”.

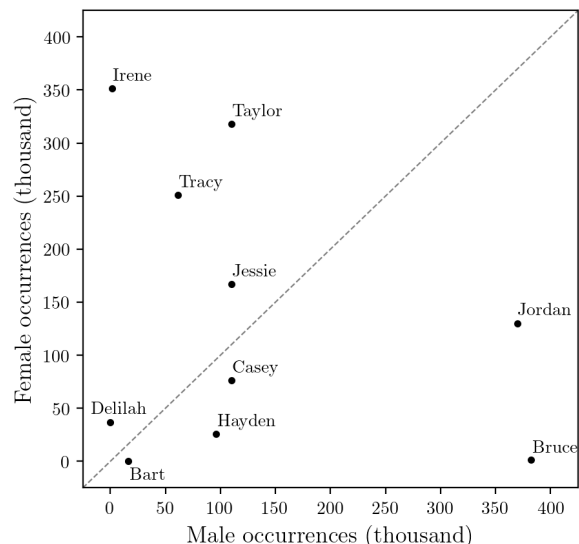


Figure 2: Frequency and gender-specificity of names in the SSA dataset

could in fact reinforce certain biases instead of mitigate them. Consider the sentence ‘Tom ... He is a successful and powerful executive’. Since *he* and *Tom* corefer, the counterfactual corpus copy will not replace *he* with *she* in this instance, and as the method involves a duplication of text, this would result in a stronger, not weaker, association between *he* and gender-stereotypic concepts present like *executive*. Even under CDS, this would still mean that biased associations are left untreated (albeit at least not reinforced). Treating names should in contrast effect a real neutralisation of bias, with the added bonus that grammaticality is maintained without the need for coreference resolution.

The United States Social Security Administration (SSA) dataset contains a list of all first names from Social Security card applications for births in the United States after 1879, along with their gender.² Figure 2 plots a few example names according to their male and female occurrences, and shows that names have varying degrees of gender-specificity.³

We fixedly associate pairs of names for swapping, thus expanding Lu et al.’s short list of gender pairs vastly. Clearly both name frequency and the degree of gender-specificity are relevant to this bipartite matching. If only frequency were considered, a more gender-neutral name (e.g. *Taylor*) could be paired with a very gender-specific name

²<https://www.ssa.gov/oact/babynames/background.html>

³The dotted line represents gender-neutrality, and more frequent names are located further from the origin.

(e.g. *John*), which would negate the gender intervention in many cases (namely whenever a male occurrence of *Taylor* is transformed into *John*, which would also result in incorrect pronouns, if present). If, on the other hand, only the degree of gender-specificity were considered, we would see frequent names (like *James*) being paired with far less frequent names (like *Sybil*), which would distort the overall frequency distribution of names. This might also result in the retention of a gender signal: for instance, swapping a highly frequent male name with a rare female name might simply make the rare female name behave as a new link between masculine contexts (instead of the original male name), as it rarely appears in female contexts.

Figure 3 shows a plot of various names’ number of primary gender⁴ occurrences against their secondary gender occurrences, with red dots for primary-male and blue crosses for primary-female names.⁵ The problem of finding name-pairs thus decomposes into a Euclidean-distance bipartite matching problem, which can be solved using the Hungarian method (Kuhn, 1955). We compute pairs for the most frequent 2500 names of each gender in the SSA dataset. There is also the problem that many names are also common nouns (e.g. *Ambler*, *Rose*, or *Mark*), which we solve using Named Entity Recognition.

4 Experimental Setup

We compare eight variations of the mitigation methods. **CDA** is our reimplementation of Lu et al.’s (2018) naïve intervention, **gCDA** uses their grammar intervention, and **nCDA** uses our new Names Intervention. **gCDS** and **nCDS** are variants of the grammar and Names Intervention using CDS. **WED40** is our reimplementation of Bolukbasi et al.’s (2016) method, which (like the original) uses a single component to define the gender subspace, accounting for > 40% of variance. As this is much lower than in the original paper (where it was 60%, reproduced in Figure 4), we define a second space, **WED70**, which uses a 2D subspace accounting for > 70% of variance. To test whether WED profits from additional names, we use the 5000 paired names in the names gazetteer as

⁴Defined as its most frequently occurring gender.

⁵The hatched area demarcates an area of the graph where no names can exist: if any name did then its primary and secondary gender would be reversed and it would belong to the alternate set.




Figure 3: Bipartite matching of names by frequency and gender-specificity

additional equalise pairs (**nWED70**).⁶ As control, we also evaluate the unmitigated space (**none**).

We perform an empirical comparison of these bias mitigation techniques on two corpora, the Annotated English Gigaword (Napoles et al., 2012) and Wikipedia. Wikipedia is of particular interest, since though its Neutral Point of View (NPOV) policy⁷ predicates that all content should be presented without bias, women are nonetheless less likely to be deemed “notable” than men of equal stature (Reagle and Rhue, 2011), and there are differences in the choice of language used to describe them (Bamman and Smith, 2014; Graells-Garrido et al., 2015). We use the annotation native to the Annotated English Gigaword, and process Wikipedia with CoreNLP (statistical coreference; bidirectional tagger). Embeddings are created using Word2Vec⁸. We use the original complex lexical input (gender-word pairs and the like) for each algorithm as we assume that this benefits each algorithm most. Expanding the set of gender-specific words for WED (following Bolukbasi et al., using a linear classifier) on Gigaword resulted in 2141 such words, 7146 for Wikipedia.⁹

⁶We use the 70% variant as preliminary experimentation showed that it was superior to WED40.

⁷https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view

⁸A CBOW model was trained over five epochs to produce 300 dimensional embeddings. Words were lowercased, punctuation other than underscores and hyphens removed, and tokens with fewer than ten occurrences were discarded.

⁹We modify or remove some phrases from the training data not included in the vocabulary of our embeddings.

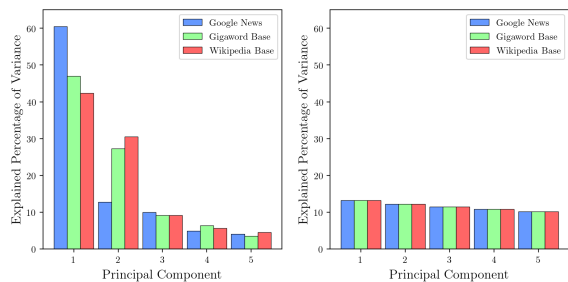


Figure 4: Variance explained by the top Principal Components of the definitional word pairs (left) and random unit vectors (right)

In our experiments, we test the degree to which the spaces are successful at mitigating direct and indirect bias, as well as the degree to which they can still be used in two NLP tasks standardly performed with embeddings, word similarity and sentiment classification. We also introduce one further, novel task, which is designed to quantify how well the embedding spaces capture an understanding of gender using non-biased analogies. Our evaluation matrix and methodology is expanded below.

Direct bias Caliskan et al. (2017) introduce the Word Embedding Association Test (WEAT), which provides results analogous to earlier psychological work by Greenwald et al. (1998) by measuring the difference in relative similarity between two sets of target words X and Y and two sets of attribute words A and B . We compute Cohen’s d (a measure of the difference in relative similarity of the word sets within each embedding; higher is more biased), and a one-sided p -value which indicates whether the bias detected by WEAT within each embedding is significant (the best outcome being that no such bias is detectable). We do this for three tests proposed by Nosek et al. (2002) which measure the strength of various gender stereotypes: art–maths, arts–sciences, and careers–family.¹⁰

Indirect bias To demonstrate indirect gender bias we adapt a pair of methods proposed by Gonen and Goldberg (2019). First, we test whether the most-biased words prior to bias mitigation remain clustered following bias mitigation. To do this, we define a new subspace, \vec{b}_{test} , using the 23 word pairs used in the Google Analogy family test subset (Mikolov et al., 2013) following Bolukbasi et al.’s (2016) method, and determine

¹⁰In the careers–family test the gender dimension is expressed by female and male first names, unlike in the other sets, where pronouns and typical gendered words are used.

the 1000 most biased words in each corpus (the 500 words most similar to \vec{b}_{test} and $-\vec{b}_{\text{test}}$) in the unmitigated embedding. For each debiased embedding we then project these words into 2D space with tSNE (van der Maaten and Hinton, 2008), compute clusters with k -means, and calculate the clusters’ V-measure (Rosenberg and Hirschberg, 2007). Low values of cluster purity indicate that biased words are less clustered following bias mitigation.

Second, we test whether a classifier can be trained to reclassify the gender of debiased words. If it succeeds, this would indicate that bias-information still remains in the embedding. We trained an RBF-kernel SVM classifier on a random sample of 1000 out of the 5000 most biased words from each corpus using \vec{b}_{test} (500 from each gender), then report the classifier’s accuracy when reclassifying the remaining 4000 words.

Word similarity The quality of a space is traditionally measured by how well it replicates human judgements of word similarity. The SimLex-999 dataset (Hill et al., 2015) provides a ground-truth measure of similarity produced by 500 native English speakers.¹¹ Similarity scores in an embedding are computed as the cosine angle between word-vector pairs, and Spearman correlation between embedding and human judgements are reported. We measure correlative significance at $\alpha = 0.01$.

Sentiment classification Following Le and Mikolov (2014), we use a standard sentiment classification task to quantify the downstream performance of the embedding spaces when they are used as a pretrained word embedding input (Lau and Baldwin, 2016) to Doc2Vec on the Stanford Large Movie Review dataset. The classification is performed by an SVM classifier using the document embeddings as features, trained on 40,000 labelled reviews and tested on the remaining 10,000 documents, reported as error percentage.

Non-biased gender analogies When proposing WED, Bolukbasi et al. (2016) use human raters to class gender-analogies as either *biased* (*woman:housewife :: man:shopkeeper*) or *appropriate* (*woman:grandmother :: man::grandfather*), and postulate that whilst biased analogies are undesirable, appropriate ones should remain. Our new analogy test uses the 506 analogies in the *fam-*

¹¹It explicitly quantifies similarity rather than association or relatedness; pairs of entities like *coffee* and *cup* have a low rating.

Method	Art–Maths		Arts–Sciences		Career–Family	
	d	p	d	p	d	p
Gigaword						
none	1.32	$< 10^{-2}$	1.50	$< 10^{-3}$	1.74	$< 10^{-4}$
CDA	0.67	.10	1.05	.02	1.79	$< 10^{-4}$
gCDA	1.16	.01	1.46	$< 10^{-2}$	1.77	$< 10^{-4}$
nCDA	-0.49	.83	0.34	.27	1.45	$< 10^{-3}$
gCDS	0.96	.03	1.31	$< 10^{-2}$	1.78	$< 10^{-4}$
nCDS	-0.19	.63	0.48	.19	1.45	$< 10^{-3}$
WED40	-0.73	.92	0.31	.28	1.24	$< 10^{-2}$
WED70	-0.73	.92	0.30	.29	1.15	$< 10^{-2}$
nWED70	0.30	.47	0.54	.19	0.59	.15
Wikipedia						
none	1.64	$< 10^{-3}$	1.51	$< 10^{-3}$	1.88	$< 10^{-4}$
CDA	1.58	$< 10^{-3}$	1.66	$< 10^{-4}$	1.87	$< 10^{-4}$
gCDA	1.52	$< 10^{-3}$	1.57	$< 10^{-3}$	1.84	$< 10^{-4}$
nCDA	1.06	.02	1.54	$< 10^{-4}$	1.65	$< 10^{-4}$
gCDS	1.45	$< 10^{-3}$	1.53	$< 10^{-3}$	1.87	$< 10^{-4}$
nCDS	1.05	.02	1.37	$< 10^{-3}$	1.65	$< 10^{-4}$
WED40	1.28	$< 10^{-2}$	1.36	$< 10^{-2}$	1.81	$< 10^{-4}$
WED70	1.05	.02	1.24	$< 10^{-2}$	1.67	$< 10^{-3}$
nWED70	-0.46	.52	-0.42	.51	0.85	.05
<i>Nosek et al.</i>	0.82	$< 10^{-2}$	1.47	$< 10^{-24}$	0.72	$< 10^{-2}$

Table 1: Direct bias results

ily analogy subset of the Google Analogy Test set (Mikolov et al., 2013) to define many such appropriate analogies that should hold even in a debiased environment, such as *boy:girl :: nephew:niece*.¹² We use a proportional pair-based analogy test, which measures each embedding’s performance when drawing a fourth word to complete each analogy, and report error percentage.

5 Results

Direct bias Table 1 presents the d scores and WEAT one-tailed p -values, which indicate whether the difference in samples means between targets X and Y and attributes A and B is significant. We also compute a two-tailed p -value to determine whether the difference between the various sets is significant.¹³

On Wikipedia, nWED70 outperforms every other method ($p < 0.01$), and even at $\alpha = 0.1$ bias was undetectable. In all CDA/S variants, the Names Intervention performs significantly better than other intervention strategies (average d for nCDS across all tests 0.95 vs. 1.39 for the best non-names CDA/S variants). Excluding the Wikipedia careers–family test (in which the CDA and CDS

¹²The entire Google Analogy Test set contains 19,544 analogies, which are usually reported as a single result or as a pair of semantic and syntactic results.

¹³Throughout this paper, we test significance in the differences between the embeddings with a two-tailed Monte Carlo permutation test at significance interval $\alpha = 0.01$ with $r = 10,000$ permutations.




Figure 5: Most biased cluster purity results

variants are indistinguishable at $\alpha = 0.01$), the CDS variants are numerically better than their CDA counterparts in 80% of the test cases, although many of these differences are not significant.

Generally, we notice a trend of WED reducing direct gender bias slightly better than CDA/S. Impressively, WED even successfully reduces bias in the careers–family test, where gender information is captured by names, which were not in WED’s gender-equalise word-pair list for treatment.

Indirect bias Figure 5 shows the V-measures of the clusters of the most biased words in Wikipedia for each embedding. Gigaword patterns similarly (see appendix). Figure 6 shows example tSNE projections for the Gigaword embeddings (“V” refers to their V-measures; these examples were chosen as they represent the best results achieved by Bolukbasi et al.’s (2016) method, Lu et al.’s (2018) method, and our new names variant). On both corpora, the new nCDA and nCDS techniques have significantly lower purity of biased-word cluster than all other evaluated mitigation techniques (0.420 for nCDS on Gigaword, which corresponds to a reduction of purity by 58% compared to the unmitigated embedding, and 0.609 (39%) on Wikipedia). nWED70’s V-Measure is significantly higher than either of the other Names variants (reduction of 11% on Gigaword, only 1% on Wikipedia), suggesting that the success of nCDS and nCDA is not merely due to their larger list of gender-words.

Figure 7 shows the results of the second test of indirect bias, and reports the accuracy of a classifier trained to reclassify previously gender biased words on the Wikipedia embeddings (Gigaword patterns similarly).¹⁴ These results reinforce the finding of the clustering experiment: once again,

¹⁴The 95% confidence interval is calculated by a Wilson score interval, i.e., assuming a normal distribution.




Figure 6: Clustering of biased words (Gigaword)




Figure 7: Reclassification of most biased words results

Method	r_s	
	Gigaword	Wikipedia
none	0.385	0.368
CDA	0.381	0.363
gCDA	0.381	0.363
nCDA	0.380	0.365
gCDS	0.382	0.366
nCDS	0.380	0.362
WED40	0.386	0.371
WED70	0.395	0.375
nWED70	0.384	0.367

Table 2: Word similarity Results

nCDS outperforms all other methods significantly on both corpora ($p < 0.01$), although it should be noted that the successful reclassification rate remains relatively high (e.g. 88.9% on Wikipedia).

We note that nullifying indirect bias associations entirely is not necessarily the goal of debiasing, since some of these may result from causal links in the domain. For example, whilst associations between *man* and *engineer* and between *man* and *car* are each stereotypic (and thus could be considered examples of direct bias), an association between *engineer* and *car* might well have little to do with gender bias, and so should not be mitigated.

Word similarity Table 2 reports the SimLex-999 Spearman rank-order correlation coefficients r_s (all are significant, $p < 0.01$). Surprisingly, the WED40 and 70 methods outperform the unmitigated embedding, although the difference in result is small (0.386 and 0.395 vs. 0.385 on Gigaword,

0.371 and 0.367 vs. 0.368 on Wikipedia). nWED70, on the other hand, performs worse than the unmitigated embedding (0.384 vs. 0.385 on Gigaword, 0.367 vs. 0.368 on Wikipedia). CDA and CDS methods do not match the quality of the unmitigated space, but once again the difference is small. It should be noted that since SimLex-999 was produced by human raters, it will reflect the human biases these methods were designed to remove, so worse performance might result from successful bias mitigation.

Sentiment classification Figure 8 shows the sentiment classification error rates for Wikipedia (Gigaword patterns similarly). Results are somewhat inconclusive. While WED70 significantly improves the performance of the sentiment classifier from the unmitigated embedding on both corpora ($p < 0.05$), the improvement is small (never more than 1.1%). On both corpora, nothing outperforms WED70 or the Names Intervention variants.

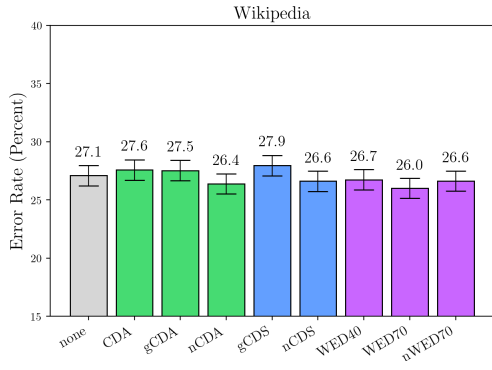


Figure 8: Sentiment classification results

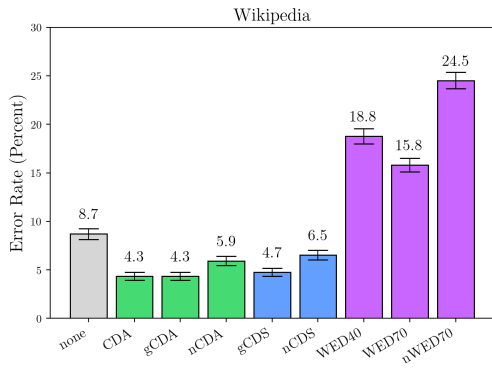


Figure 9: Non-biased gender analogy results

Non-biased gender analogies Figure 9 shows the error rates for non-biased gender analogies for Wikipedia. CDA and CDS are numerically better than the unmitigated embeddings (an effect which is always significant on Gigaword, shown in the appendices, but sometimes insignificant on Wikipedia). The WED variants, on the other hand, perform significantly worse than the unmitigated sets on both corpora (27.1 vs. 9.3% for the best WED variant on Gigaword; 18.8 vs. 8.7% on Wikipedia). WED thus seems to remove too much gender information, whilst CDA and CDS create an improved space, perhaps because they reduce the effect of stereotypical associations which were previously used incorrectly when drawing analogies.

6 Conclusion

We have replicated two state-of-the-art bias mitigation techniques, WED and CDA, on two large corpora, Wikipedia and the English Gigaword. In our empirical comparison, we found that although both methods mitigate direct gender bias and maintain the interpretability of the space, WED failed to maintain a robust representation of gender (the best variants had an error rate of 23% average when drawing non-biased analogies, suggesting that too

much gender information was removed). A new variant of CDA we propose (the Names Intervention) is the only to successfully mitigate indirect gender bias: following its application, previously biased words are significantly less clustered according to gender, with an average of 49% reduction in cluster purity when clustering the most biased words. We also proposed Counterfactual Data Substitution, which generally performed better than the CDA equivalents, was notably quicker to compute (as Word2Vec is linear in corpus size), and in theory allows for multiple intervention layers without a corpus becoming exponentially large.

A fundamental limitation of all the methods compared is their reliance on predefined lists of gender words, in particular of pairs. Lu et al.’s pairs of *manager::manageress* and *murderer::murderess* may be counterproductive, as their augmentation method perpetuates a male reading of *manager*, which has become gender-neutral over time. Other issues arise from differences in spelling (e.g. *mum* vs. *mom*) and morphology (e.g. *his* vs. *her* and *hers*). Biologically-rooted terms like *breastfeed* or *uterus* do not lend themselves to pairing either. The strict use of pairings also imposes a gender binary, and as a result non-binary identities are all but ignored in the bias mitigation literature.

Future work could extend the Names Intervention to names from other languages beyond the US-based gazetteer used here. Our method only allows for there to be an equal number of male and female names, but if this were not the case one ought to explore the possibility of a many-to-one mapping, or perhaps a probabilistic approach (though difficulties would be encountered sampling simultaneously from two distributions, frequency and gender-specificity). A mapping between nicknames (not covered by administrative sources) and formal names could be learned from a corpus for even wider coverage, possibly via the intermediary of coreference chains. Finally, given that names have been used in psychological literature as a proxy for race (e.g. Greenwald et al.), the Names Intervention could also be used to mitigate racial biases (something which, to the authors’ best knowledge, has never been attempted), but finding pairings could prove problematic. It is important that other work looks into operationalising bias beyond the subspace definition proposed by Bolukbasi et al. (2016), as it is becoming increasingly evident that gender bias is not linear in embedding space.

Acknowledgments

We would like to thank Francisco Vargas Palomo for pointing out a few typos in the proofs App. A post publication.

References

- David Bamman and Noah A. Smith. 2014. [Unsupervised discovery of biographical structure from text](#). *Transactions of the Association for Computational Linguistics*, 2:363–376.
- Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. [Man is to computer programmer as woman is to homemaker? Debiasing word embeddings](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4349–4357.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Minneapolis, Minnesota. Association for Computational Linguistics.
- Eduardo Graells-Garrido, Mounia Lalmas, and Filippo Menczer. 2015. [First women, second sex: Gender bias in Wikipedia](#). In *Proceedings of the 26th ACM Conference on Hypertext & Social Media, HT '15*, pages 165–174, New York, NY, USA. ACM.
- Anthony G. Greenwald, Debbie E. McGhee, and Jordan L. K. Schwartz. 1998. Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74(6):1464–1480.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. [SimLex-999: Evaluating semantic models with \(genuine\) similarity estimation](#). *American Journal of Computational Linguistics*, 41(4):665–695.
- Harold W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97.
- Jey Han Lau and Timothy Baldwin. 2016. [An empirical evaluation of doc2vec with practical insights into document embedding generation](#). In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 78–86, Berlin, Germany. Association for Computational Linguistics.
- Quoc Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1188–1196, Beijing, China. PMLR.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2018. [Gender bias in neural natural language processing](#). *CoRR*, abs/1807.11714.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-SNE](#). *Journal of Machine Learning Research*, 9:2579–2605.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated gigaword. In *AKBC-WEKEX Workshop at NAACL 2012*.
- Brian A. Nosek, Mahzarin R. Banaji, and Anthony G. Greenwald. 2002. Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6 1:101–115.
- Joseph Reagle and Lauren Rhue. 2011. [Gender bias in Wikipedia and Britannica](#). *International Journal of Communication*, 5(0).
- Andrew Rosenberg and Julia Hirschberg. 2007. [V-measure: A conditional entropy-based external cluster evaluation measure](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 410–420, Prague, Czech Republic. Association for Computational Linguistics.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender bias in coreference resolution: Evaluation and debiasing methods](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.

A Proofs for method from Bolukbasi et al. (2016)

We found the equations suggested in Bolukbasi et al. (2016) on the opaque side of things. So we provide here proofs missing from the original work ourselves.

Proposition 1. *Bolukbasi et al. (2016) define*

$$\vec{w} = \nu + \sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \quad (2)$$

where they define $\nu = \mu - \mu_B$. This vector is a unit vector, i.e. $\|\vec{w}\|_2 = 1$.

Proof.

$$\begin{aligned} \|\vec{w}\|_2^2 &= \vec{w}^\top \vec{w} \\ &= \left(\nu + \sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right)^\top \\ &\quad \left(\nu + \sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right) \\ &= \|\nu\|_2^2 + 2\nu^\top \left(\sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right) \\ &\quad + \left(\sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right)^\top \\ &\quad \left(\sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right) \\ &= \|\nu\|_2^2 + 2\nu^\top \left(\sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right) \\ &\quad + 1 - \|\nu\|_2^2 \\ &= 2\nu^\top \left(\sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \mu_B}{\|\vec{w}_B - \mu_B\|_2} \right) + 1 \\ &= 1 \end{aligned}$$

where we note that $\nu = \mu - \mu_B = \mu_{\perp B}$ so it is orthogonal to both \vec{w}_B and $\vec{\mu}_B$ by construction. \square

Proposition 2. *The equalise step of Bolukbasi et al. (2016) ensures that gendered pairs, e.g. man–woman, are equidistant to all gender-neutral words.*

Proof. Following Bolukbasi et al., we define \vec{e} and \vec{w} as follows:

$$\begin{aligned} \vec{e} &:= \frac{\vec{e} - \vec{e}_B}{\|\vec{e} - \vec{e}_B\|} = \frac{\vec{e}_{\perp B}}{\|\vec{e}_{\perp B}\|} \\ \vec{w} &:= \nu + \sqrt{1 - \|\nu\|_2^2} \frac{\vec{w}_B - \vec{\mu}_B}{\|\vec{w}_B - \vec{\mu}_B\|} \end{aligned}$$

Now, we have the result that

$$\vec{e} \cdot \vec{w} = \vec{e} \cdot \nu \quad (3)$$

which is the same for any \vec{e} . Now, we may compute the distance between \vec{w} and any vector \vec{e} as

$$\begin{aligned} \|\vec{e} - \vec{w}\|^2 &= (\vec{e} - \vec{w}) \cdot (\vec{e} - \vec{w}) \\ &= \vec{e} \cdot \vec{e} - 2\vec{e} \cdot \vec{w} + \vec{w} \cdot \vec{w} \\ &= 2 - 2\vec{e} \cdot \nu \end{aligned}$$

\square

B WEAT word sets

Below are listed the word sets we used for the WEAT to test direct bias, as defined by Nosek et al. (2002). Note that for the careers–family test, the target and attribute words have been reversed; that is, gender is captured by the target words, rather than the attribute words. Whilst this distinction is important in the source psychological literature (Greenwald et al., 1998), mathematically the target sets and attribute sets are indistinguishable and fully commutative.

Art–Maths Target_X: *math, algebra, geometry, calculus, equations, computation, numbers, addition*; Target_Y: *poetry, art, dance, literature, novel, symphony, drama, sculpture*; Attribute_A: *male, man, boy, brother, he, him, his, son*; Attribute_B: *female, woman, girl, sister, she, her, hers, daughter*

Arts–Sciences Target_X: *science, technology, physics, chemistry, Einstein, NASA, experiment, astronomy*; Target_Y: *poetry, art, Shakespeare, dance, literature, novel, symphony, drama*; Attribute_A: *brother, father, uncle, grandfather, son, he, his, him*; Attribute_B: *sister, mother, aunt, grandmother, daughter, she, hers, her*

Careers–Family Target_X: *John, Paul, Mike, Kevin, Steve, Greg, Jeff, Bill*; Target_Y: *Amy, Joan, Lisa, Sarah, Diana, Kate, Ann, Donna*; Attribute_A: *executive, management, professional, corporation, salary, office, business, career*; Attribute_B: *home, parents, children, family, cousins, marriage, wedding, relatives*

C Additional Gigaword results

Additional results for the Annotated English Gigaword are given here.




Figure 10: Most biased cluster purity results




Figure 11: Reclassification of most biased words results




Figure 12: Sentiment classification results




Figure 13: Non-biased gender analogy results