

Exploring the Linear Subspace Hypothesis in Gender Bias Mitigation

Francisco Vargas^δ and Ryan Cotterell^{δ,ϵ}

^δUniversity of Cambridge ^ϵETH Zürich

fav25@cam.ac.uk, rcotterell@inf.ethz.ch

Abstract

[Bolukbasi et al. \(2016\)](#) presents one of the first gender bias mitigation techniques for word embeddings. Their method takes pre-trained word embeddings as input and attempts to isolate a linear subspace that captures most of the gender bias in the embeddings. As judged by an analogical evaluation task, their method virtually eliminates gender bias in the embeddings. However, an implicit and untested assumption of their method is that the bias subspace is actually linear. In this work, we generalize their method to a kernelized, non-linear version. We take inspiration from kernel principal component analysis and derive a non-linear bias isolation technique. We discuss and overcome some of the practical drawbacks of our method for non-linear gender bias mitigation in word embeddings and analyze empirically whether the bias subspace is actually linear. Our analysis shows that gender bias is in fact well captured by a linear subspace, justifying the assumption of [Bolukbasi et al. \(2016\)](#).

1 Introduction

Pre-trained word representations are a necessity for strong performance on modern NLP tasks. These embeddings now serve as input to neural methods ([Goldberg, 2017](#)), which recently have become the standard models in the field. However, because these representations are constructed from large, human-created corpora, they naturally contain societal biases encoded in that data; gender bias is among the most well studied of these biases ([Caliskan et al., 2017](#)). Both a-contextual word embeddings ([Mikolov et al., 2013](#); [Pennington et al., 2014](#)) and contextual word embeddings ([Peters et al., 2018](#); [Devlin et al., 2019](#)) have been shown to encode gender bias ([Bolukbasi et al., 2016](#); [Caliskan et al., 2017](#); [Zhao et al., 2019](#); [May et al., 2019](#); [Karve et al., 2019](#)). More importantly, the bias in those embeddings has been shown to influence models for downstream tasks

where they are used as input, e.g. coreference resolution ([Rudinger et al., 2018](#); [Zhao et al., 2018](#)).

[Bolukbasi et al. \(2016\)](#) present one of the first methods for detecting and mitigating gender bias in word embeddings. They provide a novel linear-algebraic approach that post-processes word embeddings in order to partially remove gender bias. Under their evaluation, they find they can nearly perfectly remove bias in an analogical reasoning task. However, subsequent work ([Gonen and Goldberg, 2019](#); [Hall Maudslay et al., 2019](#)) has indicated that gender bias still lingers in the embeddings, despite [Bolukbasi et al. \(2016\)](#)'s strong experimental results. In the development of their method, [Bolukbasi et al. \(2016\)](#) make a critical and unstated assumption: Gender bias forms a linear subspace of word embedding space. In mathematics, linearity is a strong assumption and there is no reason *a-priori* why one should expect complex and nuanced societal phenomena, such as gender bias, should be represented by a linear subspace.

In this work, we present the first non-linear gender bias mitigation technique for a-contextual word embeddings. In doing so, we directly test the linearity assumption made by [Bolukbasi et al. \(2016\)](#). Our method is based on the insight that [Bolukbasi et al. \(2016\)](#)'s method bears a close resemblance to principal component analysis (PCA). Just as one can kernelize PCA ([Schölkopf et al., 1997](#)), we show that one can kernelize the method of [Bolukbasi et al. \(2016\)](#). Due to the kernelization, the bias is removed in the feature space, rather in the word embedding space. Thus, we also explore pre-image techniques ([Mika et al., 1999](#)) to project the bias-mitigated vectors back into \mathbb{R}^d .

As previously noted, there are now multiple bias removal methodologies ([Zhao et al., 2018, 2019](#); [May et al., 2019](#)) that have succeed the method by [Bolukbasi et al. \(2016\)](#). Furthermore [Gonen and Goldberg \(2019\)](#) point out multiple flaws in [Bolukbasi et al. \(2016\)](#)'s bias mitigation technique

and the aforementioned methods. Nonetheless we believe that this method has received sufficient attention from the community such that research into its properties is both interesting and useful.

We test our non-linear gender bias technique in several experiments. First, we consider the Word Embedding Association Test (WEAT; Caliskan et al., 2017); we notice that across five non-linear kernels and convex combinations thereof, there is seemingly no significant difference between the non-linear bias mitigation technique and the linear one. Secondly, we consider the professions task (Bolukbasi et al., 2016; Gonen and Goldberg, 2019) that measures how word embeddings representing different professions are potentially gender-stereotyped. Again, as with the WEAT evaluation, we find that our non-linear bias mitigation technique performs on par with the linear method. We also consider whether the non-linear gender mitigation technique removes indirect bias from the vectors (Gonen and Goldberg, 2019); yet again, we find the non-linear method performs on par with the linear methods. As a final evaluation, we evaluate whether non-linear bias mitigation hurts semantic performance. On Simlex-999 (Hill et al., 2015), we show that similarity estimates between the vectors remain on par with the linear methods. We conclude that much of the gender bias in word embeddings is indeed captured by a linear subspace, answering this paper’s titular question.

2 Bias as a Linear Subspace

The first step of Bolukbasi et al. (2016)’s technique is the discovery of a subspace $B \subset \mathbb{R}^d$ that captures most of the gender bias. Specifically, they stipulate that this space is linear. Given word embeddings that live in \mathbb{R}^d , they provide a spectral method for isolating the bias subspace. In this section, we review their approach and show how it is equivalent to principal component analysis (PCA) on a specific design (input) matrix. Then, we introduce and discuss the implicit assumption made by their work; we term this assumption the **linear subspace hypothesis** and test it in §4.

Hypothesis 1. *Gender bias in word embeddings may be represented as a linear subspace.*

2.1 Construction of a Bias Subspace

We will assume the existence of a fixed and finite vocabulary V , each element of which is a word \mathbf{w}_i . The hard-debiasing approach takes a set of N sets

$\mathcal{D} = \{D_n\}_{n=1}^N$ as input. Each set D_n contains words that are considered roughly semantically equivalent modulo their gender; Bolukbasi et al. (2016) call the D_n **defining sets**. For example, $\{man, woman\}$ and $\{he, she\}$ form two such defining sets. We identify each word with a unique integer i for the sake of our indexing notation; indeed, we exclusively reserve the index i for words. We additionally introduce the function $f : [|V|] \rightarrow [N]$ that maps an individual word to its defining set. In general, the defining sets are not limited to a cardinality of two, but in practice Bolukbasi et al. (2016) exclusively employ defining sets with a cardinality of two in their experiments. Using the sets D_n , Bolukbasi et al. (2016) define the matrix C

$$C := \sum_{n=1}^N \frac{1}{|D_n|} \sum_{i \in D_n} (\mathbf{w}_i - \boldsymbol{\mu}_{D_n})(\mathbf{w}_i - \boldsymbol{\mu}_{D_n})^\top$$

where we write \mathbf{w}_i for the i^{th} word’s embedding and the empirical mean vector $\boldsymbol{\mu}_{D_n}$ is defined as

$$\boldsymbol{\mu}_{D_n} := \frac{1}{|D_n|} \sum_{i \in D_n} \mathbf{w}_i$$

Bolukbasi et al. (2016) then extract a bias subspace B using the singular value decomposition (SVD). Specifically, they define the bias subspace to be the space spanned by the first k rows of V where

$$VSV^\top = \text{SVD}(C) \quad (1)$$

Since C is symmetric and positive semi-definite, the SVD is equivalent to an eigendecomposition as our notation in Equation 1 shows. We assume the rows of V , the eigenvectors of C , are ordered by the magnitude of their corresponding eigenvalues.

2.2 Bias Subspace Construction as PCA

As briefly noted by Bolukbasi et al. (2016), this can thus be cast as performing principal component analysis (PCA) on a recentered input matrix. We prove this assertion more formally. We first prove that the matrix C may be written as an empirical covariance matrix.

Proposition 1. *Suppose $|D_n| = 2$ for all n . Then we have*

$$C = \frac{1}{2} \sum_{i=1}^{2N} \mathbf{W}_{i,\cdot}^\top \mathbf{W}_{i,\cdot} \quad (2)$$

where we define the design matrix \mathbf{W} as:

$$\mathbf{W}_{i,\cdot} = \left(\mathbf{w}_i - \boldsymbol{\mu}_{D_{f(i)}} \right)^\top \quad (3)$$

Proof.

$$\begin{aligned}
C &= \sum_{n=1}^N \frac{1}{|D_n|} \sum_{i \in D_n} (\mathbf{w}_i - \boldsymbol{\mu}_{D_n}) (\mathbf{w}_i - \boldsymbol{\mu}_{D_n})^\top \\
&= \frac{1}{2} \sum_{n=1}^N \sum_{i \in D_n} (\mathbf{w}_i - \boldsymbol{\mu}_{D_n}) (\mathbf{w}_i - \boldsymbol{\mu}_{D_n})^\top \\
&= \frac{1}{2} \sum_{i=1}^{2N} (\mathbf{w}_i - \boldsymbol{\mu}_{D_{f(i)}}) (\mathbf{w}_i - \boldsymbol{\mu}_{D_{f(i)}})^\top \\
&= \frac{1}{2} \sum_{i=1}^{2N} \mathbf{W}_{i,\cdot}^\top \mathbf{W}_{i,\cdot}
\end{aligned}$$

where $\mathbf{W}_{i,\cdot} \in \mathbb{R}^{2N \times d}$ is defined as above. \square

Next, we show that the matrix is centered, which is a requirement for PCA.

Proposition 2. *The matrix \mathbf{W} is row-wise centered.*

Proof.

$$\begin{aligned}
\frac{1}{2} \sum_{i=1}^{2N} \mathbf{W}_{i,\cdot}^\top &= \frac{1}{2} \sum_{i=1}^{2N} (\mathbf{w}_i - \boldsymbol{\mu}_{D_{f(i)}}) \\
&= \frac{1}{2} \sum_{n=1}^N \sum_{i \in D_n} (\mathbf{w}_i - \boldsymbol{\mu}_{D_n}) \\
&= \frac{1}{2} \sum_{n=1}^N \left(-2\boldsymbol{\mu}_{D_n} + \sum_{i \in D_n} \mathbf{w}_i \right) \\
&= \frac{1}{2} \sum_{n=1}^N (2\boldsymbol{\mu}_{D_n} - 2\boldsymbol{\mu}_{D_n}) = 0
\end{aligned}$$

\square

Remark 3. *The method of Bolukbasi et al. (2016) may be considered principal component analysis performed on the matrix $2C$.*

Proof. As the algebra in proposition 1 and proposition 2 show we may formulate the problem as an SVD on a mean-centered covariance matrix. One view of PCA is performing matrix factorization on such a matrix. \square

3 Bolukbasi et al. (2016)

In this section, we review the bias mitigation technique introduced by Bolukbasi et al. (2016). When possible, we take care to reformulate their method in terms of matrix notation. They introduce a two-step process that **neutralizes** and **equalizes** the vectors to mitigate gender bias in the embeddings. The

underlying assumption of their method is that there exists a linear subspace $B \subset \mathbb{R}^d$ that captures most of the gender bias.

3.1 Neutralize

After finding the linear bias subspace B , the gist behind Bolukbasi et al. (2016)'s approach is based on elementary linear algebra. We may decompose any word vector \mathbf{w} as the sum of its orthogonal projection onto the bias subspace (range of the projection) and its orthogonal projection onto the complement of the bias subspace (null space of the projection), i.e.

$$\mathbf{w} = \mathbf{w}_B + \mathbf{w}_{\perp B} \quad (4)$$

We may then re-embed every vector as

$$\mathbf{w}_{\text{nttr}} := \mathbf{w} - \mathbf{w}_B = \mathbf{w}_{\perp B} \quad (5)$$

We may re-write this in terms of matrix notation in the following manner. Let $\{\mathbf{v}_k\}_{k=1}^K$ be an orthogonal basis for the linear bias subspace B . This may be found by taking the eigenvectors C that correspond to the top- K eigenvalues with largest magnitude. Then, we define the projection matrix onto the bias subspace as $\mathbf{P}_K = \sum_{k=1}^K \mathbf{v}_k \mathbf{v}_k^\top$ it follows that the matrix $(\mathbf{I} - \mathbf{P}_K)$ is a projection matrix on the complement of B . We can then write the neutralize step using matrices

$$\mathbf{w}_{\text{nttr}} := (\mathbf{I} - \mathbf{P}_K) \mathbf{w} \quad (6)$$

The matrix formulation of the neutralize step offers a cleaner presentation of what the neutralize step does: it projects the vectors onto the orthogonal complement of the bias subspace.

3.2 Equalize

Bolukbasi et al. (2016) decompose words into two classes. The neutral words which undergo neutralization as explained above, and the gendered words, some of which receive the equalizing treatment. Given a set of **equality sets** $\mathcal{E} = \{E_1, \dots, E_L\}$ which we can see as a greater extension of the defining sets \mathcal{D} , i.e. $E_i = \{\text{guy}, \text{gal}\}$, Bolukbasi et al. (2016) then proceed to decompose each of the words $\mathbf{w} \in \mathcal{E}$ into their gendered and neutral counterparts, setting their neutral component to a constant (the mean of the equality set) and the gendered component to its mean-centered projection into the gendered subspace:

$$\mathbf{w}_{\text{eq}} := \boldsymbol{\nu} + Z(\mathbf{w}_B - \boldsymbol{\mu}_B) \quad (7)$$

where we define the following quantities:

$$\begin{aligned} Z &:= \frac{\sqrt{1 - \|\boldsymbol{\nu}\|_2^2}}{\|\mathbf{w}_B - \boldsymbol{\mu}_B\|_2} \\ \boldsymbol{\mu} &:= \frac{1}{|E_i|} \sum_{\mathbf{w} \in E_i} \mathbf{w} \\ \boldsymbol{\nu} &:= (\mathbf{I} - \mathbf{P}_K) \boldsymbol{\mu} \end{aligned}$$

the ‘‘normalizer’’ Z ensures the vector is of unit length. This fact is left unexplained in the original work, but [Hall Maudslay et al. \(2019\)](#) provide a proof in their appendix.

4 Bias as a Non-Linear Subspace

We generalize the framework presented in [Bolukbasi et al. \(2016\)](#) and cast it to a non-linear setting by exploiting its relationship to PCA. Thus, the natural extension of [Bolukbasi et al. \(2016\)](#) is to kernelize it analogously to [Schölkopf et al. \(1997\)](#), which is the kernelized generalization of PCA. Our approach preserves all the desirable formal properties presented in the linear method of [Bolukbasi et al. \(2016\)](#).

4.1 Adapting the Design Matrix

The idea behind our non-linear bias mitigation technique is based on kernel PCA ([Schölkopf et al., 1998](#)). In short, the idea is to map the original word embeddings $\mathbf{w}_i \in \mathbb{R}^d$ to a higher-dimensional space \mathcal{H} via a function $\Phi : \mathbb{R}^d \rightarrow \mathcal{H}$. We will consider cases where \mathcal{H} is a reproducing kernel Hilbert space (RKHS) with reproducing kernel $\kappa(\mathbf{w}_i, \mathbf{w}_j) = \langle \Phi(\mathbf{w}_i), \Phi(\mathbf{w}_j) \rangle$ where the notation $\langle \cdot, \cdot \rangle$ refers to an inner product in the RKHS. Traditionally, one calls \mathcal{H} the feature space and will use this terminology throughout this work. Exploiting the reproducing kernel property, we may carry out [Bolukbasi et al. \(2016\)](#)’s bias isolation technique and construct a non-linear analogue.

We start the development of bias mitigation technique in feature space with a modification of the design matrix presented in Eq. (3). In the RKHS setting the non-linear analogue is

$$\mathbf{W}_{i,\cdot}^\Phi := \Phi(\mathbf{w}_i) - \mathbf{M}_{D_f(i)}^\Phi, \quad \forall \mathbf{w}_i \in V \quad (8)$$

where we define

$$\mathbf{M}_{D_n}^\Phi := \frac{1}{|D_n|} \sum_{i \in D_n} \Phi(\mathbf{w}_i) \quad (9)$$

As one can see, this is a relatively straightforward mapping from the set of linear operations to non-linear ones.

4.2 Kernel PCA

Using our modified design matrix, we can cast our non-linear bias mitigation technique as a form of kernel PCA. Specifically, we form the following matrix

$$\mathbf{C}^\Phi = \frac{1}{2} \sum_{i=1}^{2N} (\mathbf{W}_{i,\cdot}^\Phi)^\top \mathbf{W}_{i,\cdot}^\Phi$$

Our goal is to find the eigenvalues λ_k and their corresponding eigenfunctions $\mathbf{V}_k^\Phi \in \mathcal{H}$ by solving the eigenvalue problem

$$\mathbf{C}^\Phi \mathbf{V}_k^\Phi = \lambda_k \mathbf{V}_k^\Phi \quad (10)$$

Computing these directly from Equation 10 is impossible since \mathcal{H} ’s dimension may be prohibitively large or even infinite. However, [Schölkopf et al.](#) note that \mathbf{V}_k^Φ is spanned by $\{\Phi(\mathbf{w}_i)\}_{i=1}^{2N}$. This allows us to rewrite Eq. (10) as

$$\mathbf{V}_k^\Phi = \sum_{i=1}^{2N} \alpha_i^k \Phi(\mathbf{w}_i) \quad (11)$$

where there exist coefficients $\alpha_i^k \in \mathbb{R}$. Now, by substituting Eq. (11) and Eq. (10) into the respective terms in $\lambda \langle \Phi(\mathbf{w}_i), \mathbf{V}_k^\Phi \rangle = \langle \Phi(\mathbf{w}_i), \mathbf{C} \mathbf{V}_k^\Phi \rangle$, [Schölkopf et al. \(1997\)](#) derive a computationally feasible eigendecomposition problem. Specifically, they consider

$$2N \lambda_k \boldsymbol{\alpha}^k = \mathbf{K} \boldsymbol{\alpha}^k \quad (12)$$

where $\mathbf{K}_{ij} = \langle \Phi(\mathbf{w}_i), \Phi(\mathbf{w}_j) \rangle$. Once all the $\boldsymbol{\alpha}^k$ vectors have been estimated the inner product between an eigenfunction \mathbf{V}_k^Φ and a point \mathbf{w} can be computed as

$$\begin{aligned} \beta_k(\mathbf{w}) &= \langle \mathbf{V}_k^\Phi, \Phi(\mathbf{w}_i) \rangle \\ &= \sum_{i=1}^{2N} \alpha_i^k \kappa(\mathbf{w}_i, \mathbf{w}) \end{aligned} \quad (13)$$

A projection into the basis $\{\mathbf{V}_k^\Phi\}_{k=1}^K$ can then be carried out by applying the projection operator $\mathbf{P}_K : \mathcal{H} \rightarrow \mathcal{H}$ as follows:

$$\mathbf{P}_K \Phi(\mathbf{w}) = \sum_{k=1}^K \beta_k(\mathbf{w}) \mathbf{V}_k^\Phi \quad (14)$$

where K is the number of principal components. Projection operator \mathbf{P}_K is analogous to the linear projection \mathbf{P}_K introduced in §3.1.

4.3 Centering Kernel Matrix

We can perform the required mean-centering operations on the design matrix by centering the kernel matrix in a similar fashion to Schölkopf et al. (1998). For the case of equality sets of size 2, which is what Bolukbasi et al. use in practice, we realize that the centered design matrix reduces to pairwise differences:

$$\mathbf{W}_{i,\cdot}^{\Phi} = \frac{1}{2}(\Phi(\mathbf{w}_i) - \Phi(\mathbf{w}_j)), \quad (15)$$

$$\mathbf{w}_i, \mathbf{w}_j \in D_n \wedge \mathbf{w}_j \neq \mathbf{w}_i$$

which leads to a very simple re-centering in terms of the Gram matrices:

$$\tilde{\mathbf{K}} = \mathbf{K}^{(11)} - \mathbf{K}^{(12)} - (\mathbf{K}^{(12)})^{\top} + \mathbf{K}^{(22)}$$

where

$$\mathbf{K}_{ij}^{(xy)} = \frac{1}{2} \kappa(\mathbf{W}_{i,\cdot}^{(x)}, \mathbf{W}_{j,\cdot}^{(y)}),$$

$$\mathbf{W}_{i,\cdot}^{(z)} = \begin{cases} \mathbf{w}_{\pi_{1+i \bmod 2}(\hat{f}(\lceil \frac{i}{2} \rceil))} & z = 1 \\ \mathbf{w}_{\pi_{2-i \bmod 2}(\hat{f}(\lceil \frac{i}{2} \rceil))} & z = 2 \end{cases} \quad (16)$$

where $\hat{f} : [N] \rightarrow [|V|] \times [|V|]$ maps a defining set index to a tuple containing the word indices in the corresponding defining set and $\pi_1, \pi_2 : [|V|] \times [|V|] \rightarrow [|V|]$ are projection operators which return the first or second elements of a tuple respectively. In simpler terms, Eq. (16) is creating two matrices: matrix $\mathbf{W}^{(1)}$ which is constructed by looping over the definition sets and placing pairs within the same definition set as adjacent rows, then $\mathbf{W}^{(2)}$ is constructed in the same way but the order of the adjacent pairs is swapped relative to $\mathbf{W}^{(1)}$. Once we have this pairwise centered Gram matrix $\tilde{\mathbf{K}}$ we can apply the eigendecomposition procedure described in Equation 12 directly on $\tilde{\mathbf{K}}$. We note that carrying out this procedure using a linear kernel recovers the linear bias subspace from Bolukbasi et al. (2016).

4.4 Inner Product Correction (Neutralize)

We now focus on neutralizing and equalizing the inner products in the RKHS rather than correcting the word embeddings directly. Just as in the linear case, we can decompose the representation of a word in the RKHS into biased and neutral components

$$\Phi(\mathbf{w}) = \mathbf{P}_K^{\perp} \Phi(\mathbf{w}) + \mathbf{P}_K \Phi(\mathbf{w}), \quad (17)$$

which provides a nonlinear equivalent for Eq. (6)

$$\begin{aligned} \Phi_{\text{ntnr}}(\mathbf{w}) &= \mathbf{P}_K^{\perp} \Phi(\mathbf{w}) \\ &= \Phi(\mathbf{w}) - \mathbf{P}_K \Phi(\mathbf{w}) \end{aligned} \quad (18)$$

Proposition 4. *The corrected inner product in the feature space for two neutralized words \mathbf{z}, \mathbf{w} is given by*

$$\begin{aligned} \langle \Phi_{\text{ntnr}}(\mathbf{z}), \Phi_{\text{ntnr}}(\mathbf{w}) \rangle & \quad (19) \\ &= \kappa(\mathbf{z}, \mathbf{w}) - \sum_{k=1}^K \beta_k(\mathbf{z}) \beta_k(\mathbf{w}) \end{aligned}$$

Proof.

$$\begin{aligned} \langle \Phi_{\text{ntnr}}(\mathbf{z}), \Phi_{\text{ntnr}}(\mathbf{w}) \rangle &= \\ \langle \Phi(\mathbf{z}) - \mathbf{P}_K \Phi(\mathbf{z}), \Phi(\mathbf{w}) - \mathbf{P}_K \Phi(\mathbf{w}) \rangle & \end{aligned}$$

Applying Eq. (13) and Eq. (14)

$$\begin{aligned} &= \kappa(\mathbf{z}, \mathbf{w}) - \sum_{k=1}^K \beta_k(\mathbf{w}) \sum_{i=1}^{2N} \alpha_i^k \kappa(\mathbf{z}, \mathbf{w}_i) \\ &- \sum_{k=1}^K \beta_k(\mathbf{z}) \sum_{i=1}^{2N} \alpha_i^k \kappa(\mathbf{w}, \mathbf{w}_i) + \sum_{k=1}^K \beta_k(\mathbf{w}) \beta_k(\mathbf{z}) \\ &= \kappa(\mathbf{z}, \mathbf{w}) - \sum_{k=1}^K \beta_k(\mathbf{w}) \sum_{i=1}^{2N} \alpha_n^k \kappa(\mathbf{z}, \mathbf{w}_i) \\ &= \kappa(\mathbf{z}, \mathbf{w}) - \sum_{k=1}^K \beta_k(\mathbf{w}) \beta_k(\mathbf{z}), \end{aligned} \quad (20)$$

where

$$\beta_k(\mathbf{w}) = \mathbf{V}_k^{\Phi \top} \Phi(\mathbf{w}) = \sum_{i=1}^{2N} \alpha_i^k \kappa(\mathbf{w}, \mathbf{w}_i)$$

as derived by Schölkopf et al. (1998). \square

An advantage of this approach is that it will not rely on errors due to the approximation of the pre-image. However, it will not give us back a set of debiased embeddings. Instead, it returns a debiased metric, thus limiting the classifiers and regressors we could use. Equation Eq. (19) provides us with an $\mathcal{O}(KNd)$ approach to compute the inner product between two feature-space neutralized words.

4.5 Inner Product Correction (Equalize)

To equalize, we may naturally convert Eq. (7) to its feature-space equivalent. We define an equalizing function

$$\theta_{\text{eq}}(\mathbf{w}_i) := \nu^{\Phi} + Z^{\Phi} \mathbf{P}_K \left(\Phi(\mathbf{w}_i) - \mathbf{M}_{E_g(i)}^{\Phi} \right)$$

where we define

$$\begin{aligned}\nu^\Phi &:= P_K^\perp M_{E_{g(i)}}^\Phi \\ Z^\Phi &:= \frac{\sqrt{1 - \|\nu^\Phi\|^2}}{\|P_K(\Phi(\mathbf{w})) - M_{E_{g(i)}}^\Phi\|}\end{aligned}$$

where $g : [|V|] \rightarrow [L]$ maps an individual word index to its corresponding equality set index. Given vector dot products in the linear case follow the same geometric properties as inner products in the RKHS we can show that $\theta_{\text{eq}}(\mathbf{w})$ is unit norm follows directly from the proof for Proposition 1 in (Hall Maudslay et al., 2019) which can be found in Appendix A of Hall Maudslay et al. (2019).

Proposition 5. For any two vectors in the observed space \mathbf{w}, \mathbf{z} and their corresponding representations in feature space $\Phi(\mathbf{w}), \Phi(\mathbf{z})$ the inner product $\langle \Phi(\mathbf{w}) - P_K \Phi(\mathbf{w}), P_K \Phi(\mathbf{z}) \rangle$ is 0.

Proof.

$$\begin{aligned}\langle \Phi(\mathbf{w}) - P_K \Phi(\mathbf{w}), P_K \Phi(\mathbf{z}) \rangle &= \langle \Phi(\mathbf{w}), P_K \Phi(\mathbf{z}) \rangle - \langle P_K \Phi(\mathbf{w}), P_K \Phi(\mathbf{z}) \rangle \\ &= \sum_{k=1}^K \beta_k(\mathbf{w}) \beta_k(\mathbf{z}) - \sum_{k=1}^K \beta_k(\mathbf{w}) \beta_k(\mathbf{z}) = 0\end{aligned}$$

□

Proposition 6. For a given neutral word \mathbf{w} and a word in an equality set $e \in \mathcal{E}$ the inner product $\langle \Phi_{\text{ntnr}}(\mathbf{w}), \theta_{\text{eq}}(e) \rangle$ is invariant across members in the equality set \mathcal{E} .

Proof.

$$\begin{aligned}\langle \Phi_{\text{ntnr}}(\mathbf{w}), \theta_{\text{eq}}(e) \rangle &\stackrel{(i)}{=} \langle \Phi_{\text{ntnr}}(\mathbf{w}), \nu^\Phi \rangle \\ &= \frac{1}{|E|} \sum_{i \in E} \langle \Phi_{\text{ntnr}}(\mathbf{w}), P_K^\perp \Phi(\mathbf{w}_i) \rangle \\ &= \frac{1}{|E|} \sum_{i \in E} \langle P_K^\perp \Phi(\mathbf{w}), P_K^\perp \Phi(\mathbf{w}_i) \rangle \\ &\stackrel{(ii)}{=} \frac{1}{|E|} \sum_{i \in E} \left(\kappa(\mathbf{w}, \mathbf{w}_i) - \sum_{k=1}^K \beta_k(\mathbf{w}) \beta_k(\mathbf{w}_i) \right)\end{aligned}$$

where (i) follows from proposition 5 and (ii) follows from proposition 4. □

At this point we have completely kernelized the approach in Bolukbasi et al. (2016). Note that a linear kernel reduces to the method described in Bolukbasi et al. (2016) as we would expect. We can




Figure 1: Pre-image problem illustration for the neutralised embeddings (null-space). The plane represents the bias subspace in the RKHS.

see an initial disadvantage that equalizing via inner product correction has in comparison to Bolukbasi et al. (2016) and that is that we now require switching in between three different inner products at test time depending on whether the words are neutral or not. To overcome this in practice, we neutralize all words and do not use the equalize correction, however we present it for completeness.

5 Computing the Pre-Image

As mentioned in the previous section, a downfall of the metric correction approach is that it does not provide us embeddings that we can use in downstream tasks: the bias-mitigated embeddings only exist in feature space. Thus, when it comes to transfer tasks such as classification we are limited to kernel methods such (i.e. support vector machines). One way to resolve this problem is by obtaining the pre-image of the corrected representations in the feature space.

Finding the pre-image is a well studied problem for kernel PCA (Kwok and Tsang, 2004). The goal is to find the pre-image mappings $\Gamma : \mathcal{H}_K \oplus \mathcal{H}_K^\perp \rightarrow \mathbb{R}^d$, $\Gamma^\perp : \mathcal{H}_K^\perp \rightarrow \mathbb{R}^d$ and $\Gamma^\top : \mathcal{H}_K \rightarrow \mathbb{R}^d$ that compute (or approximate) the pre-images for $\Phi(\mathbf{w}_i)$, $\Phi_{P_K}^\perp(\mathbf{w}_i)$ and $\Phi_{P_K}(\mathbf{w}_i)$, respectively. In our case, with the pre-image mapping, the neutralize step from Bolukbasi et al. (2016) becomes

$$\mathbf{z}_i^* = \Gamma^\perp(\Phi(\mathbf{w}_i) - P_K \Phi(\mathbf{w}_i)) \quad (21)$$

In general, we will not have access to Γ^\perp so we discuss the following approximation scheme.

Additive Decomposition Approach. Alternatively, following Kandasamy and Yu (2016), we can construct an approximation to Γ that additively decomposes over the direct sum \oplus . We decompose Γ additively over the direct sum $\mathcal{H}_K \oplus \mathcal{H}_K^\perp$. That

kernel	$\kappa(\mathbf{w}_i, \mathbf{w}_j)$
Cosine	$\frac{\mathbf{w}_i^\top \mathbf{w}_j}{\ \mathbf{w}_i\ _2 \ \mathbf{w}_j\ _2}$
RBF Kernel	$\exp(-\gamma \ \mathbf{w}_i - \mathbf{w}_j\ _2^2)$
Sigmoid Kernel	$\tanh(\gamma \mathbf{w}_i^\top \mathbf{w}_j + c_0)$
Polynomial Kernel	$(\gamma \mathbf{w}_i^\top \mathbf{w}_j + c_0)^d$
Laplace Kernel	$\exp(-\gamma \ \mathbf{w}_i - \mathbf{w}_j\ _1)$
Convex Combination	$\sum_{\ell=1}^L \alpha_\ell \kappa_n(\mathbf{w}_i, \mathbf{w}_j)$,

Table 1: Different kernels used throughout experiments.

is, we assume that the pre-image mappings have the following form:

$$\Gamma(\Phi(\mathbf{w})) = \Gamma^\perp(\Phi_{P_K}^\perp(\mathbf{w})) + \Gamma^\top(\Phi_{P_K}(\mathbf{w}))$$

$$\Gamma^\perp(\Phi_{P_K}^\perp(\mathbf{w})) = \Gamma(\Phi(\mathbf{w})) - \Gamma^\top(\Phi_{P_K}(\mathbf{w}))$$

Given that we will always know that the pre-image of $\Gamma(\Phi(\mathbf{w}))$ exists and is \mathbf{w} , we can select Γ to return \mathbf{w} resulting in:

$$\Gamma^\perp(\Phi_{P_K}^\perp(\mathbf{w})) = \mathbf{w} - \Gamma^\top(\Phi^\top(\mathbf{w})) \quad (22)$$

We then learn an analytic approximation for Γ^\top using the method in [Bakır et al. \(2004\)](#). Note that most pre-imaging methods, e.g., [Mika et al. \(1999\)](#) and [Bakır et al. \(2004\)](#), are designed to approximate a pre-image Γ^\top and do not generalize to approximating the pre-image mappings Γ and Γ^\perp . This is because such methods explicitly optimise for pre-imaging representations in the space Γ^\top using points in the training set as examples of their pre-image, for the null-space Γ^\perp we have no such examples.

6 Experiments and Results

We carry out experiments across a range of benchmarks and statistical tests designed to quantify the underlying bias in word embeddings ([Gonen and Goldberg, 2019](#)). Our experiments focus on quantifying both direct and indirect bias as defined in [Gonen and Goldberg \(2019\)](#); [Hall Maudslay et al. \(2019\)](#). Furthermore we also carry out word similarity experiments using the [Hill et al. \(2015\)](#) benchmark in order to assess that our new debiased spaces still preserve the original properties of word embeddings ([Mikolov et al., 2013](#)).

6.1 Experimental Setup

Across all experiments we apply the neutralize metric correction step to all word embeddings, in con-

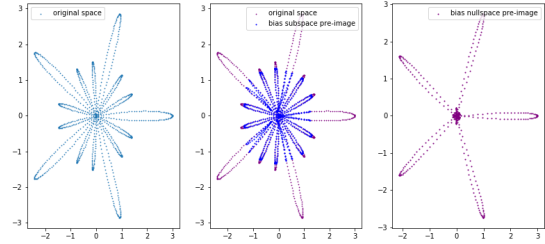


Figure 2: 2D toy example of non-linear component removal using Kernel PCA and the pre-image (neutralize step) described in §5.

trast to [Bolukbasi et al. \(2016\)](#) where the equalize step is applied to the equality sets \mathcal{E} and the neutralize step to a set of neutral words as determined in [Bolukbasi et al. \(2016\)](#). We show in [Tab. 3](#) that applying the equalize step does not bring an enhancement over neutralizing all words. We varied kernel hyper-parameters using a grid search and found that they had little effect on performance, as a result we used default initialisation strategies as suggested in [Schölkopf et al. \(1998\)](#). Unless mentioned otherwise, all experiments use the inner product correction approach introduced in §4.4.

6.2 Kernel Variations

The main kernels used throughout experiments are specified in [Tab. 1](#). We also explored the following compound kernels: (i) convex combinations of the Laplace, radial basis function (RBF), cosine and sigmoid kernels; (ii) convex combinations of cosine similarity, RBF, and sigmoid kernels; (iii) convex combinations of RBF and sigmoid kernels; (iv) polynomial kernels up to 4th degree. We only report the results on the most fundamental kernels out of the explored kernels.

6.3 Direct Bias: WEAT

The Word Embeddings Association Test [Caliskan et al. \(WEAT; 2017\)](#) is a statistical test analogous to the implicit association test (IAT) for quantifying human biases in textual data ([Greenwald and Banaji, 1995](#)). WEAT computes the difference in relative cosine similarity between two sets of target words X and Y (e.g. careers and family) and two sets of attribute words A and B (e.g. male names and female names). Formally, this quantity is Cohen’s d -measure ([Cohen, 1992](#)) also known as the effect size: The higher the measure, the more biased the embeddings. To quantify the significance of the estimated d , [Caliskan et al. \(2017\)](#) define

Targets	Original		PCA		KPCA (rbf)		KPCA (sig)		KPCPA (lap)	
	d	p	d	p	d	p	d	p	d	p
Google News Word2Vec (Mikolov et al., 2013)										
Career , Family	1.622	0.000	1.327	0.001	1.321	0.005	1.319	0.006	1.311	0.002
Math, Arts	0.998	0.017	-0.540	0.859	-0.755	0.922	-0.754	0.933	-0.024	0.507
Science , Arts	1.159	0.005	0.288	0.281	0.271	0.307	0.269	0.283	1.110	0.009
GloVe (Pennington et al., 2014)										
Career , Family	1.749	0.000	1.160	0.007	1.166	0.006	1.165	0.01	1.443	0.000
Math, Arts	1.162	0.007	0.144	0.389	0.096	0.437	0.095	0.411	0.999	0.015
Science , Arts	1.281	0.008	-1.074	0.985	-1.114	0.995	-1.112	0.993	-0.522	0.839

Table 2: WEAT results using GloVe and Google News word embeddings.

Dataset	Bolukbasi et al. (2016)		PCA	
	d	p	d	p
Google News Word2Vec (Mikolov et al., 2013)				
Career , Family	1.299	0.003	1.327	0.001
Math, Arts	-1.173	0.995	-0.540	0.859
Science , Arts	-0.509	0.832	0.288	0.281
GloVe (Pennington et al., 2014)				
Career , Family	1.160	0.000	1.160	0.007
Math, Arts	-0.632	0.887	0.144	0.389
Science , Arts	0.937	0.937	-1.074	0.985

Table 3: Effect of the equalize step

the null-hypothesis that there is no difference between the two sets of target words and the sets of attribute words in terms of their relative similarities (i.e. $d = 0$). Using this null hypothesis, Caliskan et al. (2017) then carry out a one-sided hypothesis test where failure to reject the null-hypothesis ($p > 0.05$) means that the degree of bias measured by d is not significant.

We obtain WEAT scores across different kernels (Tab. 2). We observe that the differences between the linear and the non-linear kernels is small and, in most cases, the linear kernel has a smaller value for the effect size indicating a lesser degree of bias in the corrected space. Overall, we conclude that the non-linear kernels do not reduce the linear bias as measured by WEAT further than the linear kernels. We also experiment with polynomial kernels and obtain similar results, which can be found in Tab. 7 of App. A.

Embeddings	Original	PCA	KPCA(rbf)	KPCA(sig)	KPCA(lap)
Word2Vec	0.740	0.675	0.678	0.675	0.708
Glove	0.758	0.675	0.681	0.680	0.715

Table 4: Pearson correlation coefficients of professions analogy task. All observed at significant at $\alpha = 0.05$. Indeed, all have p -values $< 10^{-30}$.

Embeddings	Original	PCA	KPCA(rbf)	KPCA(sig)	KPCA(lap)
Word2Vec	0.974	0.702	0.716	0.715	0.720
Glove	0.978	0.757	0.754	0.753	0.914

Table 5: Classification accuracy results on male versus female terms.

6.4 Professions (Gonen and Goldberg, 2019)

We consider the professions dataset introduced by Bolukbasi et al. (2016) and apply the benchmark defined in Gonen and Goldberg (2019). We find the neighbors (100 nearest neighbors) of each word using the corrected cosine similarity and count the number of male neighbors. We then report the Pearson correlation coefficient between the number of male neighbors for each word and the original bias of that word. The original bias of a word vector \mathbf{w} is given by the cosine similarity $\cos(\mathbf{w}, \mathbf{he} - \mathbf{she})$ in the original word embedding space. We can observe from the results in Tab. 4 that the non-linear kernels yield only marginally different results which in most cases seem to be slightly worse, i.e. their induced space exhibits marginally higher correlations with the original biased vector space.

6.5 Indirect Bias

Following Gonen and Goldberg (2019), we build a balanced training set of male and female words using the 5000 most biased words according to

the bias in the original embeddings as described in Section 6.4, and then train an RBF-kernel support vector machine (SVM) classifier (Pedregosa et al., 2011) on a random sample of 1000 (training set) of them to predict the gender, and evaluate its generalization on the remaining 4000 (test set). We can perform classification in our corrected RKHS with any SVM kernel $\kappa_{\text{svm}}(\Phi_{\text{nttr}}(\mathbf{w}), \Phi_{\text{nttr}}(\mathbf{z}))$ that can be written in the forms ¹ $\kappa_{\text{svm}}(\langle \mathbf{w}, \mathbf{z} \rangle)$ or $\kappa_{\text{svm}}(\|\mathbf{w} - \mathbf{z}\|^2)$ since we can use the kernel trick in our corrected RKHS $\langle \Phi_{\text{nttr}}(\mathbf{w}), \Phi_{\text{nttr}}(\mathbf{z}) \rangle = \tilde{\kappa}(\mathbf{w}, \mathbf{z})$ to compute the inputs to our SVM kernel, resulting in

$$\begin{aligned} \kappa_{\text{svm}}(\|\mathbf{w} - \mathbf{z}\|^2) \\ = \kappa_{\text{svm}}(\tilde{\kappa}(\mathbf{w}, \mathbf{w}) - 2\tilde{\kappa}(\mathbf{w}, \mathbf{z}) + \tilde{\kappa}(\mathbf{z}, \mathbf{z})). \end{aligned} \quad (23)$$

It is clear that the RBF kernel is an example of a kernel that follows Eq. (23).

We can see that the bias removal induced by non-linear kernels results in a slightly higher classification accuracy (shown in Tab. 5) of gendered words for GoogleNews Word2Vec embeddings (Mikolov et al., 2013) and a slightly lower classification accuracy for GloVe embeddings (Pennington et al., 2014) (with the exception of the Laplace kernel which has a very high classification accuracy). Overall for the RBF and the sigmoid kernels there is no improvement in comparison to the linear kernel (PCA), the Laplace kernel seems to have notably worse results than the others, still being able to classify gendered words at a high accuracy of 91.4% for GloVe embeddings.

6.6 Word Similarity: SimLex-999

The quality of a word vector space is traditionally measured by how well it replicates human judgments of word similarity. We use the SimLex-999 benchmark by Hill et al. (2015) which provides a ground-truth measure of similarity produced by 500 native English speakers. Similarity scores by our method are computed using Spearman correlation between embedding and human judgments are reported. We can observe that the metric corrections only slightly change the Spearman correlation results on SimLex-999 (Tab. 6) from the original embedding space. We can thus conclude that the embedding quality is mostly preserved.

¹Stationary kernels are sometimes written in the form $\kappa(\mathbf{w}, \mathbf{z}) = \kappa(\langle \mathbf{w}, \mathbf{z} \rangle)$ or $\kappa(\mathbf{w}, \mathbf{z}) = \kappa(\|\mathbf{w} - \mathbf{z}\|^2)$, i.e. $\kappa_{\text{RBF}}(r) = \exp(-\gamma r^2)$

Embeddings	Original	PCA	KPCA(rbf)	KPCA(sig)	KPCA(lap)
Word2Vec	0.121	0.119	0.118	0.118	0.118
Glove	0.302	0.298	0.298	0.298	0.305

Table 6: Correlation on SimLex-999 using Google-News Word2Vec and GloVe embeddings. The significance level is $\alpha = 0.05$ with $p < 0.001$.

7 Conclusion

We offer a non-linear extension to the method presented in Bolukbasi et al. (2016) by connecting its bias space construction to PCA and subsequently applying kernel PCA. We contend our extension is natural in the sense that it reduces to the method of Bolukbasi et al. (2016) in the special case when we employ a linear kernel and in the non-linear case it preserves all the desired linear properties in the feature space. This allows us to provide equivalent constructions of the neutralize and equalize steps presented.

We compare the linear bias mitigation technique to our new kernelized non-linear version across a suite of tasks and datasets. We observe that our non-linear extensions of Bolukbasi et al. (2016) show no notable performance differences across a set of benchmarks designed to quantify gender bias in word embeddings. Furthermore, the results in Tab. 7(App. A) show that gradually increasing the degree of non-linearity has again no significant change in performance for the WEAT (Caliskan et al., 2017) benchmark. Thus, we provide empirical evidence for the linear subspace hypothesis; our results suggest representing gender bias as a linear subspace is a suitable assumption. We would like to highlight that our results are specific to our proposed kernelized extensions and does not imply that all non-linear variants of (Bolukbasi et al., 2016) will yield similar results. There may very well exist a non-linear technique that works better, but we were unable to find one in this work.

References

- Gökhan H. Bakır, Jason Weston, and Bernhard Schölkopf. 2004. Learning to find pre-images. *Advances in Neural Information Processing Systems*, 16:449–456.
- 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.

- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Jacob Cohen. 1992. Statistical power analysis. *Current Directions in Psychological Science*, 1(3):98–101.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yoav Goldberg. 2017. *Neural Network Methods in Natural Language Processing*. Morgan & Claypool Publishers.
- 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 and Short Papers)*, pages 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anthony G. Greenwald and Mahzarin R. Banaji. 1995. Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1):4.
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2019. [It’s all in the name: Mitigating gender bias with name-based counterfactual data substitution](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5266–5274, Hong Kong, China. Association for Computational Linguistics.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Kirthevasan Kandasamy and Yaoliang Yu. 2016. Additive approximations in high dimensional nonparametric regression via the SALSA. In *International Conference on Machine Learning*, pages 69–78.
- Saket Karve, Lyle Ungar, and João Sedoc. 2019. Conceptor debiasing of word representations evaluated on WEAT. *arXiv preprint arXiv:1906.05993*.
- James T. Kwok and Ivor W. Tsang. 2004. The pre-image problem in kernel methods. *IEEE Transactions on Neural Networks*, 15(6):1517–1525.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. [On measuring social biases in sentence encoders](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 622–628, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sebastian Mika, Bernhard Schölkopf, Alex J. Smola, Klaus-Robert Müller, Matthias Scholz, and Gunnar Rätsch. 1999. Kernel PCA and de-noising in feature spaces. In *Advances in Neural Information Processing Systems*, pages 536–542.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. 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.
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. 1997. Kernel principal component analysis. In *International Conference on Artificial Neural Networks*, pages 583–588. Springer.
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. [Gender bias in contextualized word embeddings](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 629–634, Minneapolis, Minnesota. 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.

Targets	Original		PCA		KPCA (poly-2)		KPCA (poly-3)		KPCPA (poly-4)	
	d	p	d	p	d	p	d	p	d	p
Google News Word2Vec (Mikolov et al., 2013)										
Career , Family	1.622	0.000	1.327	0.001	1.320	0.004	1.321	0.001	1.312	0.002
Math, Arts	0.998	0.017	-0.540	0.859	-0.755	0.927	-0.755	0.933	-0.754	0.932
Science , Arts	1.159	0.005	0.288	0.281	0.271	0.312	0.272	0.305	0.272	305
GloVe (Pennington et al., 2014)										
Career , Family	1.749	0.000	1.160	0.007	1.166	0.000	1.166	0.009	1.667	0.005
Math, Arts	1.162	0.007	0.144	0.389	0.096	0.429	0.097	0.421	0.097	0.432
Science , Arts	1.281	0.008	-1.074	0.985	-1.113	0.995	-1.114	0.994	-1.114	0.992

Table 7: Results for polynomial Kernel Experiments on Glove and Google News embeddings.

A Polynomial Kernel Results

For experimental completeness, we provide direct bias experiments on WEAT using a range of polynomial kernels. The results are displayed in Tab. 7. The results for the polynomial kernels suggest the same conclusions we arrived at in the main text, i.e. a linear kernel is generally enough.