

What are the biases in my word embedding?

Nathaniel Swinger*
Lexington High School

Maria De-Arteaga*
Carnegie Mellon University

Neil Thomas Heffernan IV
Shrewsbury High School

Mark DM Leiserson
University of Maryland

Adam Tauman Kalai
Microsoft Research

December 27, 2019

Warning: This paper includes stereotypes and terms which are offensive in nature. It is important, however, for researchers to be aware of the biases contained in word embeddings.

Abstract

This paper presents an algorithm for enumerating biases in word embeddings. The algorithm exposes a large number of offensive associations related to sensitive features such as race and gender on publicly available embeddings, including a supposedly “debiased” embedding. These embedded biases are concerning in light of the widespread use of word embeddings. The associations are identified by geometric patterns in word embeddings that run parallel between people’s names and common lower-case words and phrases. The algorithm is highly unsupervised: it does not even require the sensitive groups (such as gender, race, or religion) to be pre-specified. This is desirable because: (a) it may not always be easy to identify all vulnerable groups a priori; and (b) it makes it easier to identify biases against *intersectional* groups, which depend on combinations of sensitive features. The inputs to our algorithm are a list of target tokens, e.g. names, and a word embedding. It outputs a number of Word Embedding Association Tests (WEATs) that capture various biases present in the data. We illustrate the utility of our approach on publicly available word embeddings and lists of names, and evaluate its output using crowdsourcing. We also show how removing names may not remove potential proxy bias.

1 Introduction

Bias in data representation is an important element of fairness in Artificially Intelligent systems (Barocas et al., 2017; Caliskan, Bryson, and Narayanan, 2017; Zemel et al., 2013; Dwork et al., 2012). We consider the problem of *Unsupervised Bias Enumeration* (UBE), discovering biases automatically from an unlabeled data representation. There are multiple reasons why one might want such an algorithm. First, social scientists can use it as a tool to study human bias, as data analysis is increasingly common in social studies of human biases (Garg et al., 2018; Kozłowski, Taddy, and Evans, 2018). Second, identifying bias is a natural step in “debiasing” representations (Bolukbasi et al., 2016). Finally, it can help in avoiding systems that perpetuate these biases: problematic biases can raise red flags for engineers, while little or no bias can be a useful green light indicating that a representation is usable. This awareness may be useful for identifying problems or even suggesting that a representation should not be used in a certain application. While deciding which biases are problematic is ultimately application specific, UBE may be useful in a “fair ML” pipeline.

We design a UBE algorithm for word embeddings, which are commonly used representations of tokens (e.g. words and phrases) that have been found to contain harmful bias (Bolukbasi et al., 2016). Researchers linking these biases to human biases proposed the Word Embedding Association Test (WEAT) (Caliskan, Bryson, and Narayanan, 2017). The WEAT draws its inspiration from the Implicit Association Test (IAT), a widely-used approach to measure human bias (Greenwald, McGhee, and Schwartz, 1998). An IAT $\mathcal{T} = (X_1, A_1, X_2, A_2)$ compares two sets of *target tokens* X_1 and X_2 , such as female vs. male names, and a pair of opposing sets of *attribute tokens* A_1 and A_2 , such as workplace vs. family-themed words. Average differences in a person’s response times when asked to link tokens that have anti-stereotypical vs. stereotypical relationships have been shown to indicate the strength of association between concepts. Analogously, the WEAT uses vector similarity across pairs of tokens

* Indicates equal contribution.

Word2Vec trained on Google news			fastText trained on the Web			GloVe trained on the Web		
w2v F8	w2v F11	w2v F6	fast F10	fast F7	fast F5	glove F8	glove F7	glove F5
illegal immigrant	aggravated robbery	subcontinent	n*****	jihad	s*****	turban	cartel	pornstar
drug trafficking	aggravated assault	tribesmen	f*****	militants	maid	saree	undocumented	hottie
deported	felonious assault	miscreants	dreads	caliphate	busty	hijab	culpable	nubile

Table 1: Terms associated with name groups (see Tables 3 and 6 for name groups **w2v F8**, etc.) generated from three popular pre-trained word embeddings that were rated by crowd workers as both most offensive and aligned with societal biases. These associations do *not* reflect the personal beliefs of the crowd workers or authors of this paper. See Appendix A for a discussion of the bleep-censored words.

in the sets to measure association strength. As in the case of the IAT, the inputs for a WEAT are sets of tokens \mathcal{T} predefined by researchers.

Our UBE algorithm takes as input a word embedding and a list of target tokens, and *outputs* numerous tests $\mathcal{T}_1, \mathcal{T}_2, \dots$, that are found to be statistically significant by a method we introduce for bounding false discovery rates. A crowdsourcing study of tests generated on three publicly-available word embeddings and a list of names from the Social Security Administration confirms that the biases enumerated are largely consistent with human stereotypes. The generated tests capture racial, gender, religious, and age biases, among others. Table 1 shows the name/word associations output by our algorithm that were rated most offensive by crowd workers.

Creating such tests automatically has several advantages. First, it is not feasible to manually author all possible tests of interest. Domain experts normally create such tests, and it is unreasonable to expect them to cover all possible groups, especially if they do not know which groups are represented in their data. For example, a domain expert based on the United States may not think of testing for caste discrimination, hence biases that an embedding may have against certain Indian last names may go unnoticed. Finally, if a word embedding reveals no biases, this is evidence for lack of bias. We test this by running our UBE algorithm on the supposedly debiased embedding of Bolukbasi et al. (2016).

Our approach for UBE leverages two geometric properties of word embeddings, which we call the “parallel” and “cluster” properties. The well-known parallel property is that differences between two similar token pairs, such as Mary–John and Queen–King, are often nearly parallel vectors. This suggests that among tokens in a similar topic or category, those parallel to name differences may represent biases, as was found by Bolukbasi et al. (2016) and Caliskan, Bryson, and Narayanan (2017). The cluster property, which were previously unaware of, is that the (normalized) vectors of names and words cluster into semantically meaningful groups. For names, the clusters capture social structures such as gender, religion, and others. For words, clusters of words include word categories on topics such as food, education, occupations, and sports. We use these properties to design a UBE algorithm that outputs WEATs.

Technical challenges arise around any procedure for enumerating biases. First, the combinatorial explosion of comparisons among multiple groups parallels issues in human IAT studies as aptly described by Bluemke and Friese (2008): “The evaluation of multiple target concepts such as social groups within a multi-ethnic nation (e.g. White vs. Asian Americans, White vs. African Americans, African vs. Asian Americans; Devos and Banaji, 2005) requires numerous pairwise comparisons for a complete picture”. We alleviate this problem, paralleling that work on human IATs, by generalizing the WEAT to n groups for arbitrary n . The second problem, for any UBE algorithm, is determining statistical significance to account for multiple hypothesis testing. To do this, we introduce a novel rotational null hypothesis specific to word embeddings. Third, we provide a human evaluation of the biases, contending with the difficulty that many people are unfamiliar with some groups of names.

Beyond word embeddings and IATs, related work in other subjects is worth mention. First, a body of work studies fairness properties of classification and regression algorithms (e.g. Dwork et al., 2012; Kearns et al., 2017). While our work does not concern supervised learning, it is within this work that we find one of our main motivations—the importance of accounting for intersectionality when studying algorithmic biases. In particular, Buolamwini and Gebru (2018) demonstrate accuracy disparities in image classification highlighting the fact that the magnitude of biases against an intersectional group may go unnoticed when only evaluating for each protected feature independently. Finally, while a significant portion of the empirical research on algorithmic fairness has focused on the societal biases that are most pressing in the countries where the majority of researchers currently conducting the work are based, the literature also contains examples of biases that may be of particular importance in other parts of the world (Shankar et al., 2017; Hoque et al., 2017). UBE can aspire to be useful in multiple contexts, and enable the discovery of biases in a way that relies less on enumeration by domain experts.

2 Definitions

A d -dimensional word embedding consists of a set of tokens \mathcal{W} with a nonzero vector $\mathbf{w} \in \mathbb{R}^d$ associated with each token $w \in \mathcal{W}$. Vectors are displayed in boldface. As is standard, we refer to the *similarity* between tokens v and w by the cosine of their vector angle, $\cos(\mathbf{v}, \mathbf{w})$. We write $\bar{\mathbf{v}} = \mathbf{v}/|\mathbf{v}|$ to be the vector normalized to unit-length associated with any vector $\mathbf{v} \in \mathbb{R}^d$ (or 0 if $\mathbf{v} = 0$). This enables us to conveniently write the similarity between tokens v and w as an inner product, $\cos(\mathbf{v}, \mathbf{w}) = \bar{\mathbf{v}} \cdot \bar{\mathbf{w}}$. For token set S , we write $\bar{\mathbf{S}} = \sum_{v \in S} \bar{\mathbf{v}}/|S|$ so that $\bar{\mathbf{S}} \cdot \bar{\mathbf{T}} = \text{mean}_{v \in S} \bar{\mathbf{v}} \cdot \bar{\mathbf{w}}$ is the mean similarity between pairs of tokens in sets S, T . We denote the set difference between S and T by $S \setminus T$, and we denote the first n whole numbers by $[n] = \{1, 2, \dots, n\}$.

2.1 Generalizing Word Embedding Association Tests

We assume that there is a given set of possible targets \mathcal{X} and attributes \mathcal{A} . Henceforth, since in our evaluation all targets are names and all attributes are lower-case words (or phrases), we refer to targets as names and attributes as words. Nonetheless, in principle, the algorithm can be run on any sets of target and attribute tokens. Caliskan, Bryson, and Narayanan (2017) define a WEAT statistic for two equal-sized groups of names $X_1, X_2 \subseteq \mathcal{X}$ and words $A_1, A_2 \subseteq \mathcal{A}$ which can be conveniently written in our notation as,

$$s(X_1, A_1, X_2, A_2) \stackrel{\text{def}}{=} \left(\sum_{x \in X_1} \bar{\mathbf{x}} - \sum_{x \in X_2} \bar{\mathbf{x}} \right) \cdot (\bar{\mathbf{A}}_1 - \bar{\mathbf{A}}_2).$$

In studies of human biases, the combinatorial explosion in groups can be avoided by teasing apart *Single-Category* IATs which assess associations one group at a time (e.g. Karpinski and Steinman, 2006; Penke, Eichstaedt, and Asendorpf, 2006; Bluemke and Friese, 2008). In word embeddings, we define a simple generalization for $n \geq 1$, nonempty groups X_1, \dots, X_n of arbitrary sizes and words A_1, \dots, A_n , as follows:

$$g(X_1, A_1, \dots, X_n, A_n) \stackrel{\text{def}}{=} \sum_{i=1}^n (\bar{\mathbf{X}}_i - \boldsymbol{\mu}) \cdot (\bar{\mathbf{A}}_i - \bar{\mathbf{A}})$$

where $\boldsymbol{\mu} \stackrel{\text{def}}{=} \begin{cases} \bar{\mathcal{X}} & \text{for } n = 1, \\ \sum_i \bar{\mathbf{X}}_i / n & \text{for } n \geq 2. \end{cases}$

Note that g is symmetric with respect to ordering and weights groups equally regardless of size. The definition differs for $n = 1$, otherwise $g \equiv 0$.

The following three properties motivate this as a “natural” generalization of WEAT to one or more groups.

Lemma 1. For any $X_1, X_2 \subseteq \mathcal{X}$ of equal sizes $|X_1| = |X_2|$ and any nonempty $A_1, A_2 \subseteq \mathcal{A}$,

$$s(X_1, A_1, X_2, A_2) = 2|X_1| g(X_1, A_1, X_2, A_2)$$

Lemma 2. For any nonempty sets $X \subset \mathcal{X}$, $A \subset \mathcal{A}$, let their complements sets $X^c = \mathcal{X} \setminus X$ and $A^c = \mathcal{A} \setminus A$. Then,

$$g(X, A) = 2g(X, A, X^c, A^c) = 2 \frac{|X^c|}{|\mathcal{X}|} \frac{|A^c|}{|\mathcal{A}|} g(X, A, X^c, A^c)$$

Lemma 3. For any $n > 1$ and nonempty $X_1, X_2, \dots, X_n \subseteq \mathcal{X}$ and $A_1, A_2, \dots, A_n \subseteq \bar{\mathcal{A}}$,

$$g(X_1, A_1, \dots, X_n, A_n) = \sum_{i \in [n]} g(X_i, A_i) - \sum_{i, j \in [n]} \frac{g(X_i, A_j)}{n}$$

Lemma 1 explains why we call it a generalization: for $n = 2$ and equal-sized name sets, the values are proportional with a factor that only depends on the set size. More generally, g can accommodate unequal set sizes and $n \neq 2$.

Lemma 2 shows that for $n = 1$ group, the definition is proportional the WEAT with the two groups X vs. all names \mathcal{X} and words A vs. \mathcal{A} . Equivalently, it is proportional to the WEAT between X and A and their complements.

Finally, Lemma 3 gives a *decomposition* of a WEAT into n^2 single-group WEATs $g(X_i, A_j)$. In particular, the value of a single multi-group WEAT reflects a combination of the n association strengths between X_i and A_i and n^2 disassociation strengths between X_i and A_j . As discussed on the literature on IATs, a large effect could reflect a strong association between X_1 and A_1 or X_2 and A_2 , a strong disassociation between X_1 and A_2 or X_2 and A_1 , or some combination of these factors. Proofs are deferred to Appendix B.

name	meaning	default
WE	word embedding	w2v
\mathcal{X}	set of names	SSA
n	number of target groups	12
m	number of categories	64
M	number of frequent lower-case words	30,000
t	number of words per WEAT	3
α	false discovery rate	0.05

Table 2: Inputs to the UBE algorithm.

3 Unsupervised Bias Enumeration algorithm

The inputs to our UBE algorithm are shown in Table 2. The output is m WEATs, each with n groups with associated sets of words and statistical confidences (p-values) in $[0, 1]$. Each WEAT has words from a single category, but several of the m WEATs may yield no significant associations.

At a high level, the algorithm follows a simple structure. It selects n disjoint groups of names $X_1, \dots, X_n \subset \mathcal{X}$, and m disjoint categories of lower-case words $\mathcal{A}_1, \dots, \mathcal{A}_m$. All WEATs share the same n name groups, and each WEAT has words from a single category \mathcal{A}_j , with t words associated to each X_i . Thus the WEATs can be conveniently visualized in a tabular structure.

For convenience, we normalize all word embedding vectors to be unit length. Note that we only compute cosines between them, and the cosine is simply the inner product for unit vectors. We now detail the algorithm’s steps.

3.1 Step 1: Cleaning names and defining groups

We begin with a set of names¹ \mathcal{X} , e.g., frequent first names from a database. Since word embeddings do not differentiate between words that have the same spelling but different meanings, we first “clean” the given names to remove names such as “May” and “Virginia”, whose embeddings are more reflective of other uses, such as a month or verb and a US state. Our cleaning procedure, detailed in Appendix C, is similar to that of Caliskan, Bryson, and Narayanan (2017).

We then use K-means++ clustering (from scikit-learn, Pedregosa et al., 2011, with default parameters) to cluster the normalized word vectors of the names, yielding groups $X_1 \cup \dots \cup X_n = \mathcal{X}$. Finally, we define $\mu = \sum_i \bar{\mathbf{X}}_i/n$.

3.2 Step 2: Defining word categories

To define categories, we cluster the most frequent M lower-case tokens in the word embedding into m clusters using K-means++, yielding clusters of categories $\mathcal{A}_1, \dots, \mathcal{A}_m$. The constant M is chosen to cover as many recognizable words as possible without introducing too many unrecognizable tokens. As we shall see, categories capture concepts such as occupations, food-related words, and so forth.

3.3 Step 3: Selecting words $A_{ij} \subset \mathcal{A}_j$

A test $\mathcal{T}_j = (X_1, A_{1j}, \dots, X_n, A_{nj})$ is chosen with disjoint $A_{ij} \subset \mathcal{A}_j$, each of size $t = |A_{ij}|$. To ensure disjointness,² \mathcal{A}_j is first partitioned into n “Voronoi” sets $V_{ij} \subseteq \mathcal{A}_j$ consisting of the words whose embedding is closest to each corresponding center $\bar{\mathbf{X}}_i$, i.e.,

$$V_{ij} = \left\{ w \in \mathcal{A}_j \mid i = \arg \max_{i' \in [n]} \bar{\mathbf{w}} \cdot \bar{\mathbf{X}}_{i'} \right\}$$

It then outputs A_{ij} defined as the t words maximizing the following:

$$\max_{w \in V_{ij}} (\bar{\mathbf{X}}_i - \mu) \cdot (\bar{\mathbf{w}} - \bar{\mathcal{A}}_j)$$

The more computationally-demanding step is to compute, using Monte Carlo sampling, the n p-values for \mathcal{T}_j , as described next.

¹While the set of names is an input to our system, they could also be extracted from the Embedding itself.

²If multiplicities are desired, the Voronoi sets V_{ij} could be omitted, optimizing $A_{ij} \subset \mathcal{A}_j$ directly.

3.4 Step 4: Computing p-values and ordering

To test whether the associations we find are larger than one would find if there was no relationship between the names X_i and words \mathcal{A} , we consider the following “**rotational null hypothesis**”: the words in the Embedding are generated through some process in which the alignment between names and words is random. This is formalized by imagining that a random rotation was applied (multiplying by a uniformly Haar random orthogonal matrix U) to the word embeddings but not to the name embeddings.

Specifically, to compute p-value p_{ij} for each (X_i, A_{ij}) , we first compute a score $\sigma_{ij} = (\bar{X}_i - \mu) \cdot (\bar{A}_{ij} - \bar{\mathcal{A}})$. We then compute $R = 10,000$ uniformly random orthogonal rotations $U_1, \dots, U_R \in \mathbb{R}^{d \times d}$, drawn according to the Haar measure. For each rotation, we simulate running our algorithm as if the name embeddings were transformed by U (while the word embeddings remain as is). For each rotation U_r , the sets A_{ijr} chosen to maximize $(\bar{X}_i U_r - \mu U_r) \cdot (\bar{w} - \bar{\mathcal{A}}_j)$, and the corresponding V_{ijr} and the resulting σ_{ijr} are computed. Finally, p_{ij} is the fraction of rotations for which the score $\sigma_{ijr} \geq \sigma_{ij}$ (plus an add-1 penalty standard for Monte Carlo p-values).

Furthermore, since the algorithm outputs many (hundreds) of name/word biases, the Benjamini-Hochberg (1995) procedure is used to determine a critical p-value that guarantees an α bound on the rate of false discoveries. Finally, to choose an output ordering on significant tests, the m tests are then sorted by the total scores σ_{ij} over the pairs determined significant.

4 Evaluation

To illustrate the performance of the proposed system in discovering associations, we use a database of first names provided by the Social Security Administration (SSA), which contains number of births per year by sex (F/M) (Administration, 2018). Preprocessing details are in Appendix C.

We use three publicly available word embeddings, each with $d = 300$ dimensions and millions of words: `w2v`, released in 2013 and trained on approximately 100 billion words from Google News (Mikolov et al., 2013), `fast`, trained on 600 billion words from the Web (Mikolov et al., 2018), and `glove`, also trained on the Web using the GloVe algorithm (Pennington, Socher, and Manning, 2014).

While it is possible to display the three words in each A_{ij} , the hundreds or thousands of names in each X_i cannot be displayed in the output of the algorithm. Instead, we use a simple greedy heuristic to give five “illustrative” names for each group, which are displayed in the tables in this paper and in our crowdsourcing experiments. The $k + 1^{\text{st}}$ name shown is chosen, given the first k names, so as to maximize the average similarity of the first $k + 1$ names to that of the entire set X_i . Hence, the first name is the one whose normalized vector is most central (closest to the cluster mean), the second name is the one which when averaged with the first is as central as possible, and so forth.

The WEATs can be evaluated in terms of the quality of the name groups and also their associations with words. A priori, it was not clear whether clustering name embeddings would yield any name groups or word categories of interest. For all three embeddings we find that the clustering captures latent groups defined in terms of race, age, and gender (we only have binary gender statistics), as illustrated in Table 3 for $n = 12$ clusters. While even a few clusters suffice to capture some demographic differences, more clusters yield much more fine-grained distinctions. For example, with $n = 12$ one cluster is of evidently Israeli names (see column I of table 3), which one might not consider predefining a priori since they are a small minority in the U.S. Table 6 in the Appendix shows demographic composition of clustering for other embeddings. Note that, although we do not have religious statistics for the names, several of the words in the generated associations are religious in nature suggesting religious biases as well.

Table 7 in the appendix shows the biases found in the “debiased” `w2v` embedding of Bolukbasi et al. (2016). While the name clusters still exhibit strong binary gender differences, many fewer statistically significant associations were generated for the most gender-polarized clusters.

4.1 Crowdsourcing Evaluation

We solicited ratings on the biases generated by the algorithm from US-based crowd workers on Amazon’s Mechanical Turk³ platform. The aim is to identify whether the biases found by our UBE algorithm are consistent with (problematic) biases held by society at large. To this end, we asked about society’s stereotypes, *not* personal beliefs.

We evaluated the top 12 WEATs generated by our UBE algorithm for the three embeddings, considering $n = 12$ first name groups. Our approach was simple: after familiarizing participants with the 12 groups, we showed the (statistically significant) words and name groups of a WEAT and asked them to identify which words would, stereotypically, be most associated with

³<http://mturk.com>

w2v F1	w2v F2	w2v F3	w2v F4	w2v F5	w2v F6	w2v F7	w2v F8	w2v F9	w2v F10	w2v F11	w2v F12
Amanda	Janice	Marquisha	Mia	Kayla	Kamal	Daniela	Miguel	Yael	Randall	Dashaun	Keith
Renee	Jeanette	Latisha	Keva	Carsyn	Nailah	Lucien	Deisy	Moses	Dashiell	Jamell	Gabe
Lynnea	Lenna	Tyrique	Hillary	Aislynn	Kya	Marko	Violeta	Michal	Randell	Marlon	Alfred
Zoe	Mattie	Marygrace	Penelope	Cj	Maryam	Emelie	Emilio	Shai	Jordan	Davonta	Shane
Erika	Marylynn	Takiyah	Savanna	Kaylei	Rohan	Antonia	Yareli	Yehudis	Chace	Demetrius	Stan
+581	+840	+692	+558	+890	+312	+391	+577	+120	+432	+393	+494
98% F	98% F	89% F	85% F	78% F	65% F	59% F	56% F	40% F	27% F	5% F	4% F
1983	1968	1978	1982	1993	1991	1985	1986	1989	1981	1984	1976
4% B	8% B	48% B	10% B	2% B	7% B	4% B	2% B	5% B	10% B	32% B	6% B
4% H	4% H	3% H	9% H	1% H	4% H	9% H	70% H	10% H	3% H	5% H	3% H
3% A	3% A	1% A	11% A	1% A	32% A	4% A	8% A	5% A	4% A	3% A	5% A
89% W	84% W	47% W	69% W	95% W	56% W	83% W	21% W	79% W	83% W	59% W	86% W

Table 3: Illustrative first names (greedily chosen) for $n = 12$ groups on the w2v embedding. Demographic statistics (computed a posteriori) are also shown though were not used in generation, including percentage female (at birth), mean year of birth, and percentage Black, Hispanic, Asian/Pacific Islander, and White.

Emb.	# significant	% accurate	% offensive
w2v	235	72%	35%
fast	160	80%	38%
glove	442	48%	24%

Table 4: Summary statistics for the WEATs generated using the three embeddings ($n = 12$, $m = 64$). The total number of significant name/word associations, the fraction with which the crowd’s choice of name group agreed with that of the generated WEAT (accuracy) among the top-12 WEATs, and the fraction rated as offensive.

which names group. A bonus was given for ratings that agreed with most other worker’s ratings, incentivizing workers to provide answers that they felt corresponded to widely held stereotypes.

This design was chosen over a simpler one in which WEATs are simply shown to individuals and they are asked whether or not they are stereotypical. The latter design might support confirmation bias in that people may interpret words in such a way as to confirm whatever stereotypes they are being asked about. For instance, someone may be able to justify associating the color red with almost any group, a posteriori.

Note that the task presented to the workers involved fine-grained distinctions: for each of the top-12 WEATs, at least 18 workers would each be asked to match the significant $c \leq 12$ word triples to the c name groups (each identified by five names each). For example, workers faced the triple of “registered nurse, homemaker, chairwoman” with $c = 8$ groups of names, half of which were majority female, and the most commonly chosen group matched the one generated: “Janice, Jeanette, Lenna, Mattie, Marylynn.” Across the top-12 WEATs over the three embeddings, the mean number of choices c was 8.1, yet the most commonly chosen group (plurality) agreed with the generated group 65% of the time (see Table 4). This is significantly more than one would expect from chance. The top-12 WEATs generated for w2v are shown in Table 5.

One challenge faced in this process was that, in pilot experiments, a significant fraction of the workers were not familiar with many of the names. To address this challenge, we first administered a qualification exam (common in crowdsourcing) in which each worker was shown 36 random names, 3 from each group, and was offered a bonus for each name they could correctly identify the group from which it was chosen. Only workers whose accuracy was greater than 1/2 (which happened 37% of the time) then evaluated the WEATs. Accuracy greater than 50% on a 12-way classification indicates that the groups of names were meaningful and interpretable to many workers.

Finally, we asked 13-15 workers to rate associations on a scale of 1-7 of *political incorrectness*, with 7 being “politically incorrect, possibly very offensive” and 1 being “politically correct, inoffensive, or just random.” Only those biases for which the most commonly chosen group matched the association identified by the UBE algorithm were included in this experiment. The mean ratings are shown in Table 4 and the terms present in associations deemed most offensive are presented in Table 1.

w2v F1	w2v F2	w2v F3	w2v F4	w2v F5	w2v F6	w2v F7	w2v F8	w2v F9	w2v F11	w2v F12
	cookbook, baking, baked goods	sweet potatoes, macaroni, green beans			saffron, halal, sweets	mozzarella, foie gras, caviar	tortillas, salsa, tequila	kosher, hummus, bagel	fried chicken, crawfish, grams	beef, beer, hams
herself, hers, moms	husband, homebound, grandkids	aunt, niece, grandmother	hubby, socialite, cuddle	twin sister, girls, classmate	elder brother, dowry, refugee camp			bereaved, immigrated, emigrated	younger brother, twin brother, mentally r*****	buddy, boyhood, fatherhood
hostess, cheer- leader, dietitian	registered nurse, homemaker, chairwoman		supermodel, beauty queen, stripper	helper, getter, snowboarder	shopkeeper, villager, cricketer		translator, interpreter, smuggler		cab driver, jailer, schoolboy	pitchman, retired, pundit
	log cabin, library, fairgrounds	front porch, carport, duplex	racecourse, plush, tenements	picnic tables, bleachers, concession stand	locality, mosque, slum	prefecture, chalet, sauna		synagogues, constructions, hilltop	apartment complex, barbershop, nightclub	
	parish, church, pastoral	pastor, baptized, mourners	goddess, celestial, mystical		fatwa, mosques, martyrs	monastery, papal, convent	rosary, parish priest, patron saint	rabbis, synagogue, biblical		
volleyball, gymnast, setter	athletic director, winningest coach, officiating	leading rebounder, played sparingly, incoming freshman	hooker, footy, stud	sophomore, junior, freshman	leftarm spinner, dayers, leg spinner				cornerback, tailback, wide receiver	
sorority, gymnastics, majoring	volunteer, volunteering, secretarial	guidance counselor, prekinder- garten, graduate		seventh grader, eighth grade, seniors	lecturers, institutes, syllabus		bilingual, permanent residency, occupations		incoming freshmen, schoolyard, recruiting	fulltime, professional, apprentice- ship
		civil rights, poverty stricken, nonviolent			subcontinent, tribesmen, miscreants	xenophobia, anarchist, oligarchs	leftist, drug traffickers, undocumented	disengage- ment, intifada, settlers	blacks, segregation, lynching	
tiara, blonde, sparkly	knitting, sewing, beaded	brown eyes, cream colo... wore	girly, feminine, flirty	brown hair, pair, skates	sari, turban, hijab				dreadlocks, shoulderpads, waistband	mullet, gear, helmet
					dirhams, lakhs, rupees	rubles, kronor, roulette	pesos, remittances, gross re- ceipts	shekels, settlements, corpus		
		grandjury indicted, degree murder, violating probation		child endangerment, vehicular homicide, unlawful possession	chargesheet, absconding, interrogation	absentia, tax evasion, falsification	illegal immigrant, drug trafficking, deported		aggravated robbery, aggravated assault, felonious assault	
	volunteers, crafters, baby boomers	caseworkers, evacuees, attendants	beauties, celebs, paparazzi	setters, helpers, captains	mediapersons, office bearers, newsmen				recruits, reps, sheriffs	

Table 5: The top-12 WEATs output by our UBE algorithm on the $w2v$ embedding. Columns represent name groups X_i from Table 3, rows represent categories A_j (e.g., a cluster of food-related words). Orange indicate associations where the crowd’s most commonly chosen name group agrees with that of the generated WEAT. No significant biases generated for **w2v F10**.

4.2 Potential Indirect Biases and Proxies

Naively, one may think that removing names from a dataset will remove all problematic associations. However, as suggested by Bolukbasi et al. (2016), indirect biases are likely to remain. For example, consider the `w2v` word embedding, in which *hostess* is closer to *volleyball* than to *cornerback*, while *cab driver* is closer to *cornerback* than to *volleyball*. These associations, taken from columns **F1** and **F11** of Table 5, might serve as a proxy for gender and/or race. For instance, if someone is applying for a job and their profile includes college sports words, such associations encoded in the embedding may lead to racial or gender biases in cases in which there is no professional basis for these associations. In contrast, *volunteer* being closer to *volunteers* than *recruits* may represent a definitional similarity more than a proxy, if we consider proxies to be associations that mainly have predictive power due to their correlation with a protected attribute. While defining proxies is beyond the scope of this work, we do say that $A_{ij}, A_{i'j}, A_{ij'}, A_{i'j'}$ is a *potential indirect bias* if,

$$(\overline{\mathbf{A}}_{ij} - \overline{\mathbf{A}}_{i'j}) \cdot (\overline{\mathbf{A}}_{ij'} - \overline{\mathbf{A}}_{i'j'}) > 0. \quad (1)$$

One way to interpret this definition is that if the embedding were to match the pair of word sets $\{A_{ij}, A_{i'j}\}$ to the pair of word sets $\{A_{ij'}, A_{i'j'}\}$, it would align with the way in which they were generated. For example, does the embedding predict that *hostess-cab driver* better fits *volleyball-cornerback* or *cornerback-volleyball* (but this question is asked with sets of $t = 3$ words)? Downstream, this would mean that replacing the word *cornerback* with *volleyball* on a profile would make it closer to *hostess* than *cab driver*.

We consider all possible fourtuples of significant associations, such that $1 \leq i < i' \leq n$ and $1 \leq j < j' \leq m$. In the case of `w2v`, 99% of 2,713 significant fourtuples lead to potential indirect biases according to eq. (1). This statistic is of 98% of 1,125 fourtuples and 97% of 1,796 fourtuples for the `fast` and `glove` embeddings, respectively. Hence, while names allow us to capture biases in the embedding, removing names is unlikely to be sufficient to debias the embedding.

5 Limitations

Absent clusters show the limitations of our approach and data. For example, even for large n , no clusters represent demographically significant Asian-American groups. However, in surnames (U.S. Census, Comenetz, 2016) gives a cluster “Yu, Tamashiro, Heng, Feng, Nakamura, +393” emerges that is largely Asian according to Census data (see Table 8 in Appendix D). This distinction may reflect naming practices among Asian Americans (Wu, 1999). Similarly, our approach may miss biases against smaller minorities or other groups whose names are not significantly differentiated. An example of this may be non-binary genders, although interestingly transgender terms were generated and rated as significant and consistent with human biases.

6 Conclusions and Discussion

We introduce the problem of Unsupervised Bias Enumeration. We propose and evaluate a UBE algorithm that outputs Word Embedding Association Tests. Unlike humans, where implicit tests are necessary to elicit socially unacceptable biases in a straightforward fashion, word embeddings can be directly probed to output hundreds of biases of varying natures, including numerous offensive and socially unacceptable biases.

The racist and sexist associations exposed in publicly available word embeddings raise questions about their widespread use. An important open question is how to reduce these biases.

Acknowledgments. We are grateful to Tarleton Gillespie, other colleagues and the anonymous reviewers for useful feedback.

References

- Administration, S. S. 2018. Baby names from social security card applications - national level data. Accessed 5 July 2018.
- Asim, J. 2008. *The N word: Who can say it, who shouldn't, and why*. Houghton Mifflin Harcourt.
- Barocas, S.; Crawford, K.; Shapiro, A.; and Wallach, H. 2017. The problem with bias: from allocative to representational harms in machine learning. *Special Interest Group for Computing, Information and Society (SIGCIS)*.
- Benjamini, Y., and Hochberg, Y. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57(1):289–300.

- Bianchi, C. 2014. Slurs and appropriation: An echoic account. *Journal of Pragmatics* 66:35–44.
- Bluemke, M., and Friese, M. 2008. Reliability and validity of the single-target iat (st-iat): assessing automatic affect towards multiple attitude objects. *European journal of social psychology* 38(6):977–997.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*, 4349–4357.
- Buolamwini, J., and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 77–91.
- Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334):183–186.
- Comenetz, J. 2016. Frequently occurring surnames in the 2010 census. *United States Census Bureau*.
- Devos, T., and Banaji, M. R. 2005. American= white? *Journal of personality and social psychology* 88(3):447.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, 214–226. New York, NY, USA: ACM.
- Garg, N.; Schiebinger, L.; Jurafsky, D.; and Zou, J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115(16):E3635–E3644.
- Greenwald, A. G.; McGhee, D. E.; and Schwartz, J. L. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology* 74(6):1464.
- Hoque, M.; Fatima, R. E.; Mandal, M. K.; Saquib, N.; et al. 2017. Evaluating gender portrayal in bangladeshi tv. *arXiv preprint arXiv:1711.09728*.
- Karpinski, A., and Steinman, R. B. 2006. The single category implicit association test as a measure of implicit social cognition. *Journal of personality and social psychology* 91(1):16.
- Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2017. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. *arXiv preprint arXiv:1711.05144*.
- Kozlowski, A. C.; Taddy, M.; and Evans, J. A. 2018. The geometry of culture: Analyzing meaning through word embeddings. *arXiv preprint arXiv:1803.09288*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Mikolov, T.; Grave, E.; Bojanowski, P.; Puhersch, C.; and Joulin, A. 2018. Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Penke, L.; Eichstaedt, J.; and Asendorpf, J. B. 2006. Single-attribute implicit association tests (sa-iat) for the assessment of unipolar constructs. *Experimental Psychology* 53(4):283–291.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Shankar, S.; Halpern, Y.; Breck, E.; Atwood, J.; Wilson, J.; and Sculley, D. 2017. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint arXiv:1711.08536*.
- Wu, E. D. 1999. they call me bruce, but they won't call me bruce jones: asian american naming preferences and patterns. *Names* 47(1):21–50.
- Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning fair representations. In *International Conference on Machine Learning*, 325–333.

A Offensive Stereotypes and Derogatory Terms

The authors consulted with colleagues whether to display the offensive terms and stereotypes that emerged from the embedding using our algorithms. First, regarding derogatory terms, people we consulted found the explicit inclusion of some of these terms offensive. We are also sensitive to the fact that, even in investigating them, we are ourselves using them. The terms we bleep-censor in the tables include slurs regarding race, homosexuality, transgender, and mental ability (Bianchi, 2014). In particular, these include three variants on “the n word” (Asim, 2008), *shemale*, *faggot*, *twink*, *mentally retarded*, and *rednecks*. It is not obvious that such slurs would be generated given common naming conventions. Nonetheless, many of these terms were in groups of words that matched stereotypes indicated by crowd workers.

Of course, the associations of words and groups are also offensive, but unfortunately, it is impossible to convey the nature of these associations without presenting the words in the tables associated with the groups. In an attempt to soften the effect, we use group letters rather than illustrative names or summary statistics in our tables. While this decreases the transparency, it gives the reader a choice about whether or not to examine the associated names. Some colleagues were taken aback by an initial draft, in which names and associations were displayed in the same table, and it was noted that it that may be especially offensive to individuals whose name appeared on top of a column of offensive stereotypes. For the names, we restrict our selection of names to those that had at least 1,000 occurrences in the data so that the name would not be uniquely identified with any individual.

In addition, we considered withholding the entire tables and merely presenting the rating statistics. However, we decided that, given that our concern in the analysis is uncovering that such troubling associations are being made by these tools, it was important to be clear and unflinching about what we found, and not risk obscuring the very phenomenon in our explanation.

B Proofs of Lemmas

Proof of Lemma 1. For $n = 2$, using our \bar{X} notation and their assumption $|X_1| = |X_2|$, simple algebra shows that,

$$(\bar{X}_1 - \bar{X}_2) \cdot (\bar{A}_1 - \bar{A}_2) = \frac{1}{|\bar{X}_1|} s(X_1, A_1, X_2, A_2).$$

Since $\mu = (\bar{X}_1 + \bar{X}_2)/2$, we have that $\bar{X}_1 - \mu = (\bar{X}_1 - \bar{X}_2)/2 = -(\bar{X}_2 - \mu)$, and:

$$\begin{aligned} g(X_1, A_1, X_2, A_2) &= (\bar{X}_1 - \mu) \cdot (\bar{A}_1 - \bar{\mathcal{A}}) + (\bar{X}_2 - \mu) \cdot (\bar{A}_2 - \bar{\mathcal{A}}) \\ &= \frac{\bar{X}_1 - \bar{X}_2}{2} \cdot (\bar{A}_1 - \bar{\mathcal{A}} - (\bar{A}_2 - \bar{\mathcal{A}})) \\ &= \frac{1}{2} (\bar{X}_1 - \bar{X}_2) \cdot (\bar{A}_1 - \bar{A}_2), \end{aligned}$$

which when combined with the previous equality establishes the first equation in Lemma 1. \square

Proof of Lemma 2. Since we have shown that $(\bar{X}_1 - \bar{X}_2) \cdot (\bar{A}_1 - \bar{A}_2) = 2g(X_1, A_1, X_2, A_2)$ above, we immediately have that $g(X, A) = 2g(X, A, \mathcal{X}, \mathcal{A})$. Moreover, simple algebra shows that $g(X, A, \mathcal{X}, \mathcal{A})$ and $g(X, A, X^c, A^c)$ are proportional because $\bar{X} - \bar{\mathcal{X}} = \frac{|X^c|}{|\mathcal{X}|} (\bar{X} - \bar{X}^c)$ and similarly $\bar{A} - \bar{\mathcal{A}} = \frac{|A^c|}{|\mathcal{A}|} (\bar{A} - \bar{A}^c)$. \square

Proof of Lemma 3. Follows simply from the definition of g and μ for $n \geq 2$ and $n = 1$. \square

C Preprocessing names and words

C.1 Preprocessing first names from SSA dataset

The SSA dataset (Administration, 2018) has partial coverage for earlier years and includes all names with at least 5 births, we use only years 1938-2017 and select only the names that appeared at least 1,000 times, which cover more than 99% of the data by population. From this data, we extract the fraction of female and male births for each name as well as the mean year of birth. Of course, we select only the names appearing in the embedding.

Note that the mean of the fraction of females among our names is significantly greater than 50%, even though the US population is nearly balanced in binary gender demographics. The subtle reason is there is greater variability in female names in the data, whereas the most common names are more often male. That is, the data have fewer predominantly male first names




Figure 1: A plot of log-probability (y-axis) vs. word embedding index (x-axis) for the last name data and the word2vec word embedding. Orange points represent last names we keep and blue points are outliers we remove. As expected from Zipf’s law, the probabilities and frequencies exhibit a power-law relationship. Names removed from the data by our classifier, displayed in red, are typically words that have other more common uses than as last names.

in total with more people being given those names on average. Since we are including each name only once, this increases the female representation in the population.⁴

C.2 Preprocessing last names from U.S. Census

A dataset of last names is made publicly available by the Census Bureau of the United States and contains last names occurring at least 100 times in the 2010 census (Comenetz, 2016), broken down by percentage of race, including White, Black, Hispanic, Asian and Pacific Islander, and Native American. Again we filter for names that appear at least 1,000 times and apply the binary classification procedure described in Section 3.1 to clean the data.

C.3 “Cleaning” names

Caliskan, Bryson, and Narayanan (2017) apply a simple procedure in which they remove the 20% of words whose mean similarity to the other names is smallest. We apply a similar but slightly more sophisticated procedure by training an linear Support Vector Machine (scikit-learn’s LinearSVC, Pedregosa et al., 2011, with default parameters) to distinguish the input names from an equal number of non-names chosen randomly from the most frequent 50,000 words in the embedding. We then remove the 20% of names with smallest margin in the direction identified by the linear classifier.

Figure 1 illustrates the effect of cleaning the last names and shows that the names that tend to be removed are those that violate Zipf’s law.

⁴We performed similar experiments on a sample of names drawn according to the population and, while the names are gender balanced, the clusters exhibit less diversity and most often simply are split by gender and age – one can even have an entire cluster solely consisting of people named *Michael*.

C.4 Preprocessing words

To identify the most frequent M words in the embedding, we first restrict to tokens that consist only of the 26 lower-case English letters or spaces for embeddings that contain phrases. We also omit lower-case tokens when the upper-case version of the token is more frequent. For instance, the lower-case token “john” is removed because “John” is more frequent.

D Biases in different lists/embeddings

Table 6 shows the names from other embeddings. Table 7 shows the biases found in the “debiased” w_{2v} embedding of Bolukbasi et al. (2016), while Table 8 show last-name biases generated from the w_{2v} embeddings.

fast F1	fast F2	fast F3	fast F4	fast F5	fast F6	fast F7	fast F8	fast F9	fast F10	fast F11	fast F12
Nakesha Keisha	Carolyn Nichole	Tamara Emi	Lillian Lucinda	Alejandra Maricella	Katelyn Jayda	Ahmed Shanti	Landon Keenan	Stephan Nahum	Marquell Antwan	Greg Willie	Gerardo Renato
Kandyce Kamilah	Mel Tawnya	Isabella Karina	Velda Antoinette	Ona Fabiola	Shalyn Jaylyn	Mariyah Siddharth	Skye Courtland	Sabastian Philippe	Dakari Pernell	Edward Jefferey	Pedro Genaro
Rachal +702	Deirdre +821	Joli +622	Flossie +478	Sulema +400	Evie +851	Yasmin +288	Luke +576	Jarek +312	Jarred +440	Russ +474	Matteo +234
98% F 1980	98% F 1972	97% F 1987	96% F 1972	93% F 1984	90% F 1993	64% F 1992	22% F 1991	9% F 1987	6% F 1984	4% F 1973	2% F 1987
29% B 3% H 1% A 66% W	4% B 2% H 2% A 91% W	5% B 9% H 6% A 80% W	14% B 9% H 6% A 71% W	2% B 64% H 8% A 25% W	3% B 2% H 2% A 93% W	6% B 4% H 33% A 56% W	5% B 1% H 3% A 90% W	6% B 9% H 4% A 80% W	34% B 3% H 2% A 61% W	8% B 3% H 5% A 84% W	1% B 65% H 7% A 27% W

glove F1	glove F2	glove F3	glove F4	glove F5	glove F6	glove F7	glove F8	glove F9	glove F10	glove F11	glove F12
Elsie Carlotta	Brenda Katie	Claudia Tiara	Patrica Caren	Kylee Shaye	Laticia Jayci	Alejandra Epifanio	Amina Yair	Eldridge Tad	Damion Ronney	Kevin Ernest	Gustavo Etienne
Elizabeth Dovie	Janette Liza	Lena Melina	Mikala Cherise	Tayla Latasha	Shalanda Kalynn	Monalisa Eulalia	Rani Danial	Godfrey Asa	Winford Tavaris	Haley Matt	Lorenzo Emil
Gladys +263	Debra +396	Sasha +359	Lorine +889	Jessi +520	Noelani +1270	Alicea +395	Safa +396	Renard +434	Tylor +627	Gilbert +429	Roberto +218
99% F 1972	98% F 1974	95% F 1987	94% F 1973	89% F 1987	83% F 1978	68% F 1985	58% F 1989	18% F 1979	11% F 1982	7% F 1979	6% F 1987
15% B 11% H 6% A 68% W	4% B 3% H 3% A 89% W	6% B 12% H 7% A 73% W	7% B 3% H 2% A 88% W	9% B 3% H 3% A 85% W	14% B 28% H 2% A 55% W	1% B 67% H 9% A 22% W	5% B 4% H 22% A 68% W	13% B 3% H 4% A 80% W	11% B 2% H 2% A 84% W	7% B 3% H 4% A 85% W	3% B 41% H 6% A 50% W

deb. F1	deb. F2	deb. F3	deb. F4	deb. F5	deb. F6	deb. F7	deb. F8	deb. F9	deb. F10	deb. F11	deb. F12
Denise Audrey	Kayla Lynae	Evelyn Marquetta	Marquisha Madalynn	Zoe Nana	Kamal Nailah	Nicolas Carmella	Luis Deisy	Michal Astrid	Shaneka Dondre	Randall Scarlett	Brian Ernie
Maryalice Sonja	Gabe Tayla	Gaylen Gaye	Celene Nyasia	Crystal Georgiana	Kalan Aisha	Adrien Stefania	Alexandro Elsa	Ezra Armen	Laquanda Tavon	Windell Corrin	Matthew Kennedy
Glenna +714	Staci +845	Eula +506	Lanora +819	Sariyah +512	Rony +334	Raphael +322	Eliazar +538	Juliane +282	Tanesha +688	Coley +407	Wayne +313
99% F 1971	81% F 1989	80% F 1969	78% F 1984	71% F 1984	62% F 1991	59% F 1984	56% F 1986	54% F 1987	49% F 1983	29% F 1982	5% F 1974
4% B 3% H 3% A 89% W	4% B 3% H 2% A 91% W	17% B 6% H 4% A 72% W	5% B 3% H 3% A 89% W	10% B 9% H 11% A 70% W	6% B 5% H 32% A 56% W	6% B 5% H 5% A 73% W	1% B 16% H 8% A 18% W	2% B 6% H 3% A 88% W	49% B 3% H 2% A 45% W	9% B 3% H 4% A 83% W	5% B 3% H 5% A 87% W

w2v L1	w2v L2	w2v L3	w2v L4	w2v L5	w2v L6	w2v L7	w2v L8	w2v L9	w2v L10	w2v L11	w2v L12
Moser Persson	Stein Zucker	Boyer Lasher	Romano Klimas	Murphy Nagle	Cantrell Wooddell	Gauthier Medeiros	Burgess Willson	Gaines Derouen	Lal Haddad	Mendez Aguillon	Yu Tamashiro
Pagel Runkel	Avakian Sobel	Sawin Stoudt	Pecoraro Arnone	Igoe Crosbie	Maness Newcomb	Lafrance Lounsbury	Hatton Mutch	Gaskins Aubrey	Mensah Vora	Aispuro Forero	Heng Feng
Wagner +3035	Tepper +775	Mcintire +3013	Morreale +1416	Dillon +665	Greathouse +2444	Renard +756	Patten +2818	Rodgers +1779	Omer +423	Jurado +1913	Nakamura +393
1% B 2% H 1% A 94% W	2% B 3% H 1% A 93% W	3% B 2% H 1% A 92% W	1% B 6% H 1% A 91% W	4% B 3% H 1% A 90% W	8% B 2% H 1% A 86% W	8% B 4% H 1% A 85% W	12% B 3% H 1% A 81% W	34% B 3% H 1% A 60% W	15% B 7% H 1% A 46% W	1% B 80% H 5% A 12% W	1% B 3% H 79% A 11% W

Table 6: The first name clusters from the *fast*, *glove* and *debiased* embeddings, followed by last name clusters from the *w2v* embedding. Demographic statistics (computed a posteriori) are also shown though were not used in generation, including percentage female (at birth), mean year of birth, and percentage Black, Hispanic, Asian/Pacific Islander, and White.

deb. F1	deb. F2	deb. F3	deb. F4	deb. F5	deb. F6	deb. F7	deb. F8	deb. F9	deb. F10
professor emeritus, registered nurse, adjunct professor	eighth grader, seventh grader, sixth grader	lifelong resident, postmaster, homemaker	granddaughter, grandson, daughter	bloke, chap, hubby	shopkeeper, villager, elder brother	mobster, chef, restaurateur	translator, interpreter, notary	mathematician, physicist, researcher	cousin, jailer, roommate
volunteering, homebound, nurse practitioner	seniors, eighth grade, boys	grandparents, aunts, elderly	graduated, grandchildren, siblings	bedtime, marital, bisexual	expatriate, hostels, postgraduate		undocumented, farmworkers, bilingual		blacks, academically, mentally r*****
	medley, solo, trio	bluegrass, bandleader, banjo	trombone, percussionist, clarinet		artiste, verse, remix	maestro, accordion, operas	flamenco, tango, vibes	avant garde, violinist, techno	rapper, gospel, hip hop
	volleyball, softball, roping	bass fishing, rodeo, deer hunting	wearing helmet, horseback riding, snorkeling	racecourse, footy, footballing	cricket, badminton, cricketing	peloton, anti doping, gondola		luge, biathlon, chess	basketball, sprints, lifting weights
		rural, fairgrounds, tract	westbound, southbound, eastbound	foreshore, tenements, tourist attraction	slum, headquarter, minarets	seaside, boutiques, countryside	barangays, squatters, plazas	settlements, prefecture, inhabitants	
		supper, barbecue, chili	macaroni, green beans, pancakes		halal, sweets, hummus	pizzeria, mozzarella, pasta	tortillas, salsa, tequila	kosher, vodka, bagel	
					dirhams, emirate, riyals	euros, francs, vintages	peso, reais, nationalized	supervisory board, zloty, ruble	
		pastor, church, parish	baptized, sisters, brothers	mystical, witch, afterlife	fatwa, mosque, martyrs	nuns, papal, monastery		rabbis, synagogue, commune	
	captains, bridesmaids, grads	caretakers, grandmothers, superintendents	cousins, helpers, friends	punters, blokes, celebs	mediapersons, office bearers, shopkeepers				rappers, recruits, officers
			clan, overthrow, starvation		subcontinent, rulers, tribals		leftist, indigenous peoples, peasants	rightist, disengagement, oligarchs	civil rights, segregation, racial
					rupees, dinars, crores		pesos, remittances, cooperatives	shekels, rubles, kronor	
		convicted felon, felony convictions, probate	child endangerment, unlawful possession, vehicular homicide		chargesheet, absconding, petitioner	absentia, annulment, penitentiary			aggravated robbery, aggravated assault, felonious assault

Table 7: The top-12 WEATs output by our UBE algorithm on the “debiased” w_{2v} embedding of Bolukbasi et al. (2016), again with $n = 12$. Despite being debiased, demographic statistics (again computed a posteriori) reveal names still cluster by gender, but the extreme gender clusters have many fewer statistically significant associations. For instance, the most male groups **deb. F11** and **deb. F12** are not shown because no significant associations were generated.

w2v L1	w2v L2	w2v L3	w2v L4	w2v L5	w2v L6	w2v L7	w2v L8	w2v L9	w2v L10	w2v L11	w2v L12	
potato salad, pretzels, chocolate cake	kosher, bagel, hummus	pumpkin, brownies, donuts	mozzarella, pasta, deli	pint, whiskey, cheddar	pecans, grits, watermelon	maple syrup, foie gras	cider, lager, malt	fried chicken, crawfish, sweet potatoes	sweets, saffron, mango	tortillas, salsa, tequila	noodles, dumplings, soy sauce	
concentration, camp, extermination, postwar	disengagement, neocons, intifada			unionists, sectarian, pedophiles			province, separatist, sovereignty	antisocial behavior, cricket, asylum seekers	blacks, segregation, civil rights	non governmental, miscreants, encroachments	drug traffickers, leftist, undocumented	hyun, bian, motherland
	co founder, venture capitalist, psychotherapist	assessor, wildlife biologist, secretary treasurer	restaurant, plumber, fire-fighter	solicitor, selector, handicapper	jailer, rancher, appraiser		schoolboy, barrister, chap	cheerleader, bailiff, recruiter	shopkeeper, aspirant, taxi driver	translator, smuggler, interpreter	villager, vice, housewife	
	synagogues, skyscraper, studio	log cabin, zoning ordinance, barn	pizzeria, borough, firehouse	pubs, racecourse, western suburbs	fair-grounds, acre tract, concession stand	rink, cottage, chalet	disused, derelict, leisure		locality, slum, hostel		prefecture, guesthouse, metropolis	
	authors, hedgefund managers, creators	crafters, hobbyists, racers	mobsters, restaurateurs, captains	gardai, lads, footballers	sheriffs, folks, appraisers	skaters, premiers, mushers	blokes, householders, solicitors		mediapersons, newsmen, office bearers		migrant workers, maids, civil servants	
	rabbis, synagogue, biblical		papal, pontiff, convent	archdiocese, clerical, diocese	denomination, pastor, church		vicar, creationism, traditionalists	pulpit, preaching, preach	fatwa, fasting, sufferings	rosary, parish priest, patron saint	commune, monks, temples	
	shekels, settlements, nonprofit	mill levy, assessed valuation, tax abatement			millage, payday lenders, appropriations		unfair dismissal, attendances, takings		rupees, lakhs, dirhams	pesos, remittances, indigent	baht, overseas, income earners	
	pollster, liberal, moderates				commissioners, countywide, statewide	ridings, selectmen, byelection		desegregation, uncommitted, voter registration	panchayat, candidature, localities	barangay, immigration reform, congresswoman	plenary session, landslide, multiracial	
	insider trading, attorneys, lawsuit	felonious assault, drug paraphernalia, criminal mischief			sheriff, meth lab, jailers	impaired driving, criminal negligence, penitentiary	affray, bailiffs, aggravated burglary	aggravated robbery, racially charged, probation violation	absconding, charge-sheet, complainant	illegal immigrant, drug trafficking, deadly weapon		
						loonie, francs, takeovers	sharemarket, credit crunch, gilts		load shedding, microfinance, rupee	peso, reais, nationalization	cross strait, yuan, ringgit	
walleye, lakes, aquarium	transatlantic, iceberg, flotilla				crappie, bass fishing, boat ramp			shad, barrier islands, grouper	mangroves, jetty, kite	sardines, tuna, archipelago	mainland, seaweed, island	
feedlot, barley, wheat		cornfield, pumpkins, alfalfa			mowing, deer hunting, pasture				agro, saplings, livelihood	farmworkers, coca, sugarcane	bamboo, cassava, palm oil	

Table 8: The top-12 WEATs output by our UBE algorithm on the w2v embedding for *last names*. The corresponding name groups are presented in Table 6.