

Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting

Maria De-Arteaga¹, Alexey Romanov², Hanna Wallach³, Jennifer Chayes³, Christian Borgs³,
Alexandra Chouldechova¹, Sahin Geyik⁴, Krishnaram Kenthapadi⁴, Adam Tauman Kalai³

¹Carnegie Mellon University, ²University of Massachusetts Lowell, ³Microsoft Research, ⁴LinkedIn

ABSTRACT

We present a large-scale study of gender bias in occupation classification, a task where the use of machine learning may lead to negative outcomes on peoples' lives. We analyze the potential allocation harms that can result from semantic representation bias. To do so, we study the impact on occupation classification of including explicit gender indicators—such as first names and pronouns—in different semantic representations of online biographies. Additionally, we quantify the bias that remains when these indicators are “scrubbed,” and describe proxy behavior that occurs in the absence of explicit gender indicators. As we demonstrate, differences in true positive rates between genders are correlated with existing gender imbalances in occupations, which may compound these imbalances.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Applied computing** → **Document management and text processing**.

KEYWORDS

Supervised learning, algorithmic fairness, gender bias, online recruiting, automated hiring, compounding injustices.

ACM Reference Format:

Maria De-Arteaga¹, Alexey Romanov², Hanna Wallach³, Jennifer Chayes³, Christian Borgs³, Alexandra Chouldechova¹, Sahin Geyik⁴, Krishnaram Kenthapadi⁴, Adam Tauman Kalai³, ¹Carnegie Mellon University, ²University of Massachusetts Lowell, ³Microsoft Research, ⁴LinkedIn. 2019. Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting. In *FAT* '19: Conference on Fairness, Accountability, and Transparency, January 29–31, 2019, Atlanta, GA, USA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3287560.3287572>

1 INTRODUCTION

The presence of automated decision-making systems in our daily lives is growing. As a result these systems play an increasingly active role in shaping our future. Far from being passive players that consume information, automated decision-making systems are participating actors: their predictions today affect the world we live in tomorrow. In particular, they determine many aspects of how

we experience the world, from the news we read and the products we shop for to the job postings we see. The increased prevalence of machine learning has therefore been accompanied by a growing concern regarding the circumstances and mechanisms by which such systems may reproduce and augment the various forms of discrimination and injustices that are present in today's society.

One domain in which the use of machine learning is increasingly popular—and in which unfair practices can lead to particularly negative consequences—is that of online recruiting and automated hiring. Maintaining an online professional presence has become increasingly important for people's careers, and this information is often used as input to automated decision-making systems that advertise open positions and recruit candidates for jobs and other professional opportunities. In order to perform these tasks, a system must be able to accurately assess people's current occupations, skills, interests, and “potential.” However, even the simplest of these tasks—determining someone's current occupation—can be non-trivial. Although this information may be provided in a structured form on some professional networking platforms, this is not always the case. As a result, recruiters often browse candidates' websites in an attempt to manually determine their current occupations. Machine learning promises to reduce this burden; however, as we will explain in this paper, occupation classification is susceptible to gender bias, stemming from existing gender imbalances in occupations.

To study gender bias in occupation classification, we created a new dataset of hundreds of thousands of online biographies, written in English, from the Common Crawl corpus. Because biographies are typically written in the third person by their subjects (or people familiar with their subjects) and because pronouns are gendered in English, we were able to extract (likely) self-identified binary gender from the biographies. We note, though, that this binary model is a simplification that fails to capture important aspects of gender and erases people who do not fit within its assumptions.

Using this dataset, we predicted people's occupations by performing multi-class classification using three different semantic representations: bag-of-words, word embeddings, and deep recurrent neural networks. For each representation, we considered two scenarios: (1) where explicit gender indicators are available to the classifier, (2) where explicit gender indicators are “scrubbed” to promote fairness or to comply with regulations or laws. We define explicit gender indicators to be information, such as first names and gendered pronouns, that make it possible to determine gender. We note that the practice of “scrubbing” explicit gender indicators and other sensitive attributes is not unique to machine learning, and is often used as a way to mitigate the effects of implicit and explicit bias on decisions made by humans. For example, gender diversity

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAT '19, January 29–31, 2019, Atlanta, GA, USA*

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6125-5/19/01...\$15.00

<https://doi.org/10.1145/3287560.3287572>

in orchestras was significantly improved by the introduction of “blind” auditions, where candidates play behind a curtain [24].

To quantify gender bias, we compute the true positive rate (TPR) gender gap—i.e., the difference in TPRs between genders—for each occupation. The TPR for a given gender and occupation is defined as the proportion of people with that gender and occupation that are correctly predicted as having that occupation. We also compute the correlation between these TPR gender gaps and existing gender imbalances in occupations, and show how this may compound these imbalances; we connect this finding with an existing notion of indirect discrimination in political philosophy. We show that “scrubbing” explicit gender indicators reduces the TPR gender gaps, while maintaining overall classifier accuracy. However, we also show that significant TPR gender gaps remain in the absence of explicit gender indicators, and that these gaps are correlated with existing gender imbalances. For orchestra auditions, the sounds made by candidates’ shoes mean that a curtain is not sufficient to make an audition “blind.” It is therefore common practice to additionally roll out a carpet or to ask candidates to remove their shoes [24]. By analogy, “scrubbing” explicit gender indicators is like introducing a curtain—the sounds made by the candidates’ shoes remain.

Our paper has two main takeaways: First, “scrubbing” explicit gender indicators is not sufficient to remove gender bias from an occupation classifier. Second, even in the absence of such indicators, TPR gender gaps are correlated with existing gender imbalances in occupations; occupation classifiers may therefore compound existing gender imbalances. Although we focus on gender bias, we note that other biases, such as those involving race or socioeconomic status, may also be present in occupation classification or in other tasks related to online recruiting and automated hiring. We structure our analysis so as to inform discussions about these biases as well.

In the next section, we provide a brief overview of related work. We then describe our data collection process in Section 3 and outline our methodology in Section 4, before presenting our analysis and results in Section 5. We conclude with a discussion in Section 6.

2 RELATED WORK

Recent work has studied the ways in which stereotypes and other human biases may be reflected in semantic representations such as word embeddings [10, 13, 21]. Natural language processing researchers have also studied gender bias in coreference resolution [34, 44], showing that systems perform better when linking a gender pronoun to an occupation in which that gender is over-represented than to an occupation in which it is underrepresented. Gender bias has also been studied in YouTube’s autocaptioning [38], where researchers found a higher word error rate for female speakers. In the context of language identification, researchers have also investigated racial bias, showing that African-American English is often misclassified as non-English [8]. Finally, machine learning methods for identifying toxic comments exhibit disproportionately high false positive rates for words like *gay* and *homosexual* [17].

In the context of structured data, there have been extensive discussions about proxy behavior that may occur when sensitive attributes are not explicitly available but can be determined from other attributes [4, 33, 43]. Related discussions have focused on the phenomenon of differential subgroup validity [3], where the

choice of attributes may disadvantage groups for whom the chosen attributes are not equally predictive of the target label [12]. Barocas and Selbst [4] discussed these issues in the context of automated hiring; Kim [28] explained how data-driven decisions that systematically bias people’s access to opportunities relate to existing antidiscrimination legislation, identifying voids that may need to be filled to account for potential risks stemming from automated decision-making systems. Researchers have also discussed making available sensitive attributes as a means to improve fairness [18], as well as various ways to use these attributes [20, 33]. Finally, although our paper does not directly consider ranking scenarios, fairness in ranking is particularly relevant to discussions about gender bias in online recruiting and automated hiring [7, 14, 22, 40, 42].

We quantify gender bias by computing the TPR gender gap—i.e., the difference in TPRs between genders—for each occupation. This notion of bias is closely related to the equality of opportunity fairness metric of Hardt et al. [25]. We choose to focus on TPR gender gaps because they enable us to study the ways in which gender imbalances may be compounded; in turn, we relate this to compounding injustices [26]—an existing notion of indirect discrimination in political philosophy that holds that it is a general moral duty to refrain from taking actions that would harm people when those actions are informed by, and would compound, prior injustices suffered by those people. We show that the TPR gender gaps are correlated with existing gender imbalances in occupations. As a result, occupation classifiers compound injustices when existing gender imbalances are attributable to historical discrimination.

Our paper is also closely related to research on gender bias in hiring [5, 23, 35, 36]. In particular, Bertrand and Mullainathan [6] conducted an experiment in which they responded to help-wanted ads using fictitious resumes, varying names so as to signal gender and race, while keeping everything else the same. They were therefore able to measure the effect of (inferred) gender and race on the likelihood of being called for an interview. Similarly, we study the effect of explicit gender indicators on occupation classification.

Computational linguistics researchers have explored the use of lexical and syntactic features to infer authors’ genders [15, 30]. Given that our dataset consists of online biographies, our paper is also related to research on differences between the ways that men and women represent themselves. In the context of online professional presences, Altenburger et al. [2] analyzed self-promotion in LinkedIn, finding that women are more modest than men in expressing accomplishments and are less likely to use free-form fields. Researchers have also studied differences in volubility between men and women [11], showing that women’s fear of being highly voluble is justified by the fact that both men and women negatively evaluate highly voluble women. Moving beyond self-representation, Niven and Zilber [32] analyzed congressional websites and found that differences between the ways that the media portray men and women in Congress cannot be explained by differences between the ways that they portray themselves. Meanwhile, Smith et al. [37] analyzed attributes used to describe men and women in performance evaluations, showing that negative attributes are more often used to describe women than men. This research on representation by others relates to our paper because we cannot be sure that the online biographies in our dataset were actually written by their subjects.

3 DATA COLLECTION PROCESS

To study gender bias in occupation classification, we created a new dataset using the Common Crawl. Specifically, we identified online biographies, written in English, by filtering for lines that began with a name-like pattern (i.e., a sequence of two capitalized words) followed by the string “is a(n) (xxx) title,” where *title* is an occupation from the BLS Standard Occupation Classification system.¹ We identified the twenty-eight most frequent occupations based on their appearance in a small subset of the Common Crawl. In a few cases, we merged occupations. For example, we created the occupation *professor* by merging occupations that consist of *professor* and a modifier, such as *economics professor*. Having identified the most frequent occupations, we processed WET² files from sixteen distinct crawls from 2014 to 2018, extracting online biographies corresponding to those occupations only. Finally, we performed de-duplication by treating biographies as duplicates if they had the same first name, last name, and occupation, and either no middle name was present or one middle name was a prefix of the other. The resulting dataset consists of 397,340 biographies spanning twenty-eight different occupations. Of these occupations, *professor* is the most frequent, with 118,400 biographies, while *rapper* is the least frequent, with 1,406 biographies (see Figure 1). The longest biography is 194 tokens, while the shortest is eighteen; the median biography length is seventy-two tokens. We note that the demographics of online biographies’ subjects differ from those of the overall workforce, and that our dataset does not contain all biographies on the Internet; however, neither of these factors is likely to undermine our findings.

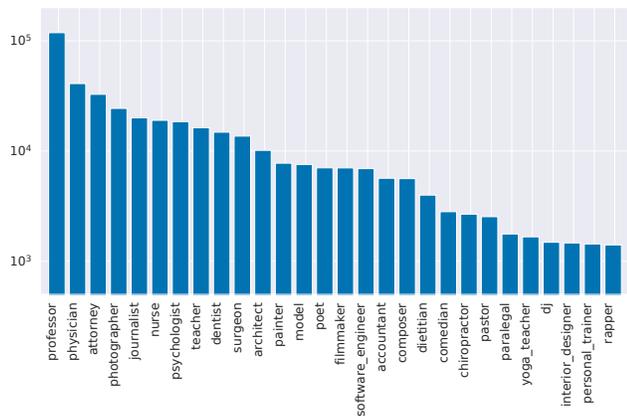


Figure 1: Distribution of the number of biographies for the twenty-eight different occupations, shown on a log scale.

Because some occupations have a high gender imbalance, our validation and testing splits must be large enough that every gender and occupation are sufficiently represented. We therefore used stratified-by-occupation splits, with 65% of the biographies (258,370) designated for training, 10% (39,635 biographies) designated for validation, and 25% (99,335 biographies) designated for testing.

A complete implementation that reproduces the dataset can be found in the source code available at <http://aka.ms/biasbios>.

¹<https://www.bls.gov/soc/>

²WET is a special file format containing cleaned text extracted from webpages.

4 METHODOLOGY

We used our dataset to predict people’s occupations, taken from the first sentence of their biographies as described in the previous section, given the remainder of their biographies. For example, consider the hypothetical biography *Nancy Lee is a registered nurse. She graduated from Lehigh University, with honours in 1998. Nancy has years of experience in weight loss surgery, patient support, education, and diabetes.* The goal is to predict *nurse* from *She graduated from Lehigh University, with honours in 1998. Nancy has years of experience in weight loss surgery, patient support, education, and diabetes.*

We used three different semantic representations of varying complexity: bag-of-words (BOW), word embeddings (WE), and deep recurrent neural networks (DNN). When using the BOW and WE representations, we used a one-versus-all logistic regression as the occupation classifier; to construct the DNN representation, we started with word embeddings as input and then trained a DNN to predict occupations in an end-to-end fashion. For each representation, we considered two scenarios: (1) where explicit gender indicators—e.g., first names and pronouns—are available to the classifier, (2) where explicit gender indicators are “scrubbed.” For example, these scenarios correspond to predicting the occupation *nurse* from the text *[She] graduated from Lehigh University, with honours in 1998. [Nancy] has years of experience in weight loss surgery, patient support, education, and diabetes*, with and without the bracketed words.

4.1 Semantic Representations

Bag-of-words. The BOW representation encodes the i^{th} biography as a sparse vector x_i^{BOW} . Each element of this vector corresponds to a word type in the vocabulary, equal to 1 if the biography contains a token of this type and 0 otherwise. Despite recent successes of using more complex semantic representations for document classification, the BOW representation provides a good baseline and is still widely used, especially in scenarios where interpretability is important. To predict occupations, we trained a one-versus-all logistic regression with L_2 regularization using our dataset’s training split represented using the BOW representation.

Word embeddings. The WE representation encodes the i^{th} biography as a vector x_i^{WE} , obtained by averaging the fastText word embeddings [9, 31] for the word types present in that biography.³ The WE representation is surprisingly effective at capturing non-trivial semantic information [1]. To predict occupations, we trained a one-versus-all logistic regression with L_2 regularization using our dataset’s training split represented using the WE representation.

Deep recurrent neural networks. To construct the DNN representation, we started with the fastText word embeddings as input and then trained a DNN to predict occupations in an end-to-end fashion. We used an architecture similar to that of Yang et al. [41], but with just one bi-directional recurrent neural network at the level of words and with gated recurrent units (GRUs) [16] instead of long short-term memory units; this model uses an attention mechanism—an integral part of modern neural network architectures [39]. Our choice of architecture was motivated by a desire to use a relatively simple model that would be easy to interpret.

³We note that the fastText word embeddings were trained using the Common Crawl, albeit using a different subset than the one we used to create our dataset.

Formally, given the i^{th} biography represented as a sequence of tokens w_i^1, \dots, w_i^T , we start by replacing each token w_i^t with the fastText word embedding for that word type to yield e_i^1, \dots, e_i^T . The DNN then uses a GRU to process the biography in both forward and reverse directions and concatenates the corresponding hidden states from both directions to re-represent the t^{th} token as follows:

$$\vec{h}_i^t = \overrightarrow{\text{GRU}}(e_i^t, h_i^{t-1}) \quad (1)$$

$$\overleftarrow{h}_i^t = \overleftarrow{\text{GRU}}(e_i^t, h_i^{t+1}) \quad (2)$$

$$h_i^t = [\vec{h}_i^t; \overleftarrow{h}_i^t]. \quad (3)$$

Next, the DNN projects each hidden state h_i^t to the attention dimension k_a via a fully connected layer with weights W_a and b_a , and transforms the result into an unnormalized scalar u_i^t via a vector w_a :

$$\hat{u}_i^t = \tanh(W_a h_i^t + b_a) \quad (4)$$

$$u_i^t = w_a^T \hat{u}_i^t. \quad (5)$$

Each scalar is then normalized to yield an attention weight:

$$\alpha_i^t = \frac{\exp(u_i^t)}{\sum_{t'=1}^T \exp(u_i^{t'})}. \quad (6)$$

Finally, we obtain the DNN representation via a weighted sum:

$$x_i^{\text{DNN}} = \sum_{t=1}^T \alpha_i^t h_i^t. \quad (7)$$

The DNN makes predictions as follows:

$$\hat{y}_i = \text{softmax}(W_0 x_i^{\text{DNN}} + b_0), \quad (8)$$

where \hat{y}_i is the predicted occupation for the i^{th} biography.

We trained the DNN using our dataset’s training split and a standard cross-entropy loss applied to the output of the last layer.

4.2 Explicit Gender Indicators

For each semantic representation, we considered two scenarios. In the first scenario, the representation included all word types, meaning that explicit gender indicators are available to the occupation classifier. In the second scenario, we “scrubbed” explicit gender indicators prior to creating the representation, meaning that these indicators are not available to the occupation classifier. Specifically, we deleted the subject’s first name, along with the words *he, she, her, his, him, hers, himself, herself, mr, mrs,* and *ms* from each biography.

5 ANALYSIS AND RESULTS

In this section, we analyze the potential allocation harms that can result from semantic representation bias. To do this, we study the performance of the occupation classifier for each semantic representation, with and without explicit gender indicators, as described in the previous section. The classifiers’ overall accuracies are shown in Figure 2. We start by analyzing gender bias for the scenario in which the semantic representations include all word types, including explicit gender indicators. We then analyze gender bias in the scenario in which explicit gender indicators are “scrubbed,” and use the DNN’s per-token attention weights to understand proxy behavior that occurs in the absence of explicit gender indicators.

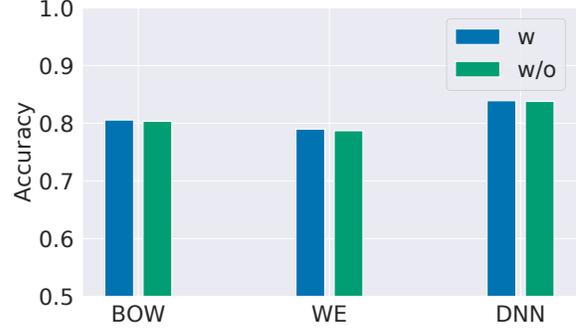


Figure 2: Occupation classifier accuracy for each semantic representation, with and without explicit gender indicators.

5.1 With Explicit Gender Indicators

True positive rate gender gap. For each semantic representation, we quantify gender bias by using our dataset’s testing split to calculate the occupation classifier’s TPR gender gap—i.e., the difference in TPRs between binary genders g and $\sim g$ —for each occupation y :

$$\text{TPR}_{g,y} = P[\hat{Y} = y | G = g, Y = y] \quad (9)$$

$$\text{Gap}_{g,y} = \text{TPR}_{g,y} - \text{TPR}_{\sim g,y}, \quad (10)$$

where \hat{Y} and Y are random variables representing the predicted and target labels (i.e., occupations) for a biography and G is a random variable representing the binary gender of the biography’s subject.

Defining the percentage of people with gender g in occupation y as $\pi_{g,y} = P[G = g | Y = y]$, Figure 3 shows $\text{Gap}_{\text{female},y}$ versus $\pi_{\text{female},y}$ for each occupation y for the BOW representation with explicit gender indicators; Figure 4 depicts the same information for all three representations, with and without explicit gender indicators.

Compounding imbalance. We define the gender imbalance of occupation y as $\frac{\pi_{g,y}}{\pi_{\sim g,y}}$; gender g is underrepresented if $\frac{\pi_{g,y}}{\pi_{\sim g,y}} < 1$ or, equivalently, if $\pi_{g,y} < 0.5$. The gender imbalance is compounded if the underrepresented gender has a lower TPR than the overrepresented gender—e.g., if $\text{Gap}_{g,y} < 0$ and g is underrepresented.

Theorem 1. If $\pi_{g,y} < 0.5$ and $\text{Gap}_{g,y} < 0$, then

$$P[G = g | Y = \hat{Y} = y] < \pi_{g,y}. \quad (11)$$

PROOF. Via Bayes theorem,

$$P[G = g | Y = \hat{Y} = y] = \frac{\pi_{g,y} \text{TPR}_{g,y}}{P[\hat{Y} = y | Y = y]}. \quad (12)$$

If $\pi_{g,y} < \pi_{\sim g,y}$ and $\text{TPR}_{g,y} < \text{TPR}_{\sim g,y}$, then

$$\frac{P[G = g | Y = \hat{Y} = y]}{P[G = \sim g | Y = \hat{Y} = y]} = \frac{\pi_{g,y} \text{TPR}_{g,y}}{\pi_{\sim g,y} \text{TPR}_{\sim g,y}} < \frac{\pi_{g,y}}{\pi_{\sim g,y}}, \quad (13)$$

so the gender imbalance for the true positives in occupation y is larger than the initial gender imbalance in that occupation. \square

As explained in Section 2, if the initial gender imbalance is due to prior injustices, an occupation classifier will compound these injustices, which may correspond to indirect discrimination [26].

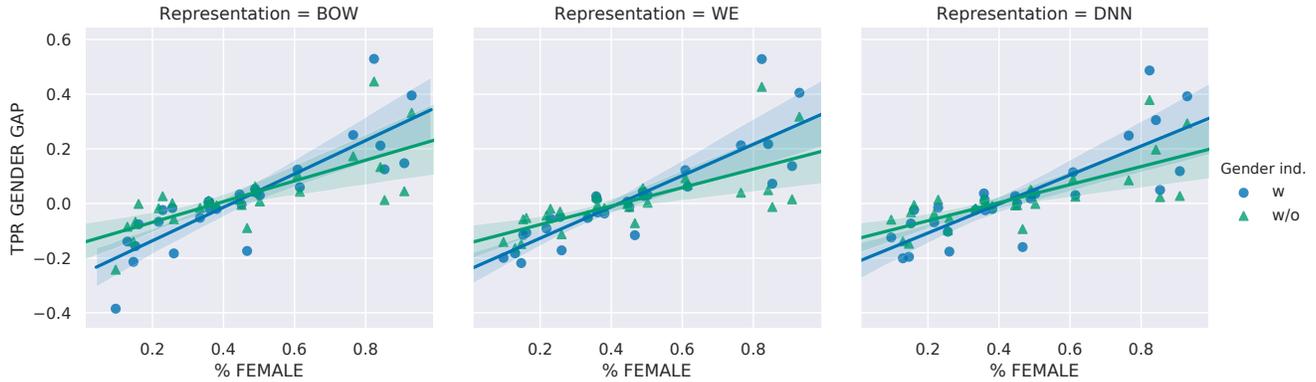


Figure 4: $\text{Gap}_{\text{female}, y}$ versus $\pi_{\text{female}, y}$ for each occupation y for all three semantic representations, with and without explicit gender indicators. Correlation coefficients: BOW-w 0.85; BOW-wo 0.74; WE-w 0.86; WE-wo 0.71; DNN-w 0.82, DNN-wo 0.74.

Table 1: Pairs of occupations with the largest values of $\Pi_{\text{male}, (y^1, y^2)}$ —i.e., the percentage of men’s biographies that are only correctly predicted as y^2 when their indicators are swapped for which the predicted label changes from y^1 .

y^1	y^2	$\Pi_{\text{male}, (y^1, y^2)}$
attorney	paralegal	7.1%
architect	interior designer	4.7%
professor	dietitian	4.3%
photographer	interior designer	3.5%
teacher	yoga teacher	3.3%

Table 2: Pairs of occupations with the largest values of $\Pi_{\text{female}, (y^1, y^2)}$ —i.e., the percentage of women’s biographies that are only correctly predicted as y^2 when their indicators are swapped for which the predicted label changes from y^1 .

y^1	y^2	$\Pi_{\text{female}, (y^1, y^2)}$
model	rapper	14.7%
teacher	pastor	8.5%
professor	software engineer	6.5%
professor	surgeon	4.8%
physician	surgeon	3.8%

representing the biography’s target label (i.e., occupation). In turn, this would mean that the TPRs for genders g and $\sim g$ are identical:

$$\text{TPR}_{g, y} = P[\hat{Y} = y | G = g, Y = y] \tag{17}$$

$$= P[\hat{Y} = y | G = \sim g, Y = y] \tag{18}$$

$$= \text{TPR}_{\sim g, y}, \tag{19}$$

where $\hat{Y} = f(\tilde{X}^R)$ is a random variable representing the predicted label (i.e., occupation) for \tilde{X}^R . Moreover, it would also mean that

$$P[G = g | \tilde{X}^R = \tilde{x}^R, Y = y] = P[G = \sim g | \tilde{X}^R = \tilde{x}^R, Y = y], \tag{20}$$

making it impossible to predict the gender of a “scrubbed” biography’s subject belonging to occupation y better than random.

In order to determine whether “scrubbing” explicit gender indicators is sufficient to remove all information about gender, we used a balanced subsample of our dataset to predict people’s gender. We created a subsampled training split by first discarding from our dataset’s training split all occupations for which there were not at least 1,000 biographies for each gender. For each of the remaining twenty-one occupations, we then subsampled 1,000 biographies for each gender to yield 42,000 biographies, balanced by occupation and gender. To create a subsampled validation split, we first identified the occupation and gender from those represented in the subsampled training split with the smallest number of biographies in our dataset’s validation split. Then, we subsampled that number of biographies from our dataset’s validation split for each of the twenty-one occupations represented in the subsampled training split and each gender. We created a subsampled testing split similarly. When using the BOW and WE representations, we used a logistic regression with L_2 regularization as the gender classifier; to construct the DNN representation, we started with word embeddings as input and then trained a DNN to predict gender in an end-to-end fashion, similar to the methodology described in Section 4.

Using the subsampled testing split, we find that the gender classifier for the BOW representation has an accuracy of 65.5%, while the DNN representation has an accuracy of 68.2%. These accuracies are higher than 50%, so “scrubbing” explicit gender indicators is not sufficient to remove all information about gender. This finding is reinforced by the scatterplot in Figure 5, which shows log frequency versus correlation with $G = \text{female}$ for each word type in the vocabulary. It is clear from this scatterplot that deleting all words that are correlated with gender would not be feasible.

True positive rate gender gap and compounding imbalance. For each semantic representation, we again quantify gender bias by using our (original) dataset’s testing split to calculate the occupation classifier’s TPR gender gap for each occupation. Figure 4 shows $\text{Gap}_{\text{female}, y}$ versus $\pi_{\text{female}, y}$ for each occupation y for all three representations, with and without explicit gender indicators.

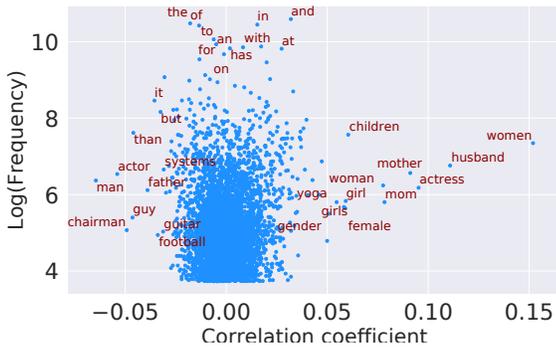


Figure 5: Scatterplot of log frequency versus correlation with $G = \text{female}$ for each word type in the vocabulary.

“Scrubbing” explicit gender indicators reduces the TPR gender gaps, while the classifiers’ accuracies (shown in Figure 2) remain roughly the same; however, for some occupations, $\text{Gap}_{\text{female},y}$ is still very large. Moreover, because there is still a positive correlation between the TPR gender gap for an occupation y and the gender imbalance in that occupation, “scrubbing” explicit gender indicators will not prevent the classifiers from compounding gender imbalances.

We note that compounding imbalances are especially problematic if people repeatedly encounter such classifiers—i.e., if an occupation classifier’s predictions determine the data used by subsequent occupation classifiers. Who is offered a job today will affect the gender (im)balance in that occupation in the future. If a classifier compounds existing gender imbalances, then the underrepresented gender will, over time, become even further underrepresented—a phenomenon sometimes referred to as the “leaky pipeline.”

To illustrate this phenomenon, we performed simulations using the DNN representation in which the candidate pool at time $t + 1$ is defined by the true positives at time t . Defining the percentage of people with gender g in occupation y at time t as $\pi_{g,y}^{(t)}$, we fit a linear regression to the TPR gender gaps for different values of $\pi_{g,y}^{(t)}$:

$$\widehat{\text{Gap}}_{g,y}^{(t)} = \pi_{g,y}^{(t)} \beta_1 + \beta_0. \quad (21)$$

Using this regression model, we are then able to calculate the percentage of people with gender g in occupation y at time $t + 1$:

$$\pi_{g,y}^{(t+1)} = \frac{\pi_{g,y}^{(t)} \text{TPR}_{g,y}^{(t)}}{\pi_{\sim g,y}^{(t)} (\text{TPR}_{g,y}^{(t)} + \widehat{\text{Gap}}_{g,y}^{(t)}) + \pi_{g,y}^{(t)} \text{TPR}_{g,y}^{(t)}}. \quad (22)$$

Figure 6 shows $\pi_{g,y}^{(t)}$ for $t = 0, \dots, 10$; each subplot corresponds to a different initial gender imbalance. Over time, the gender imbalances compound. We note that there are many different TPR pairs $\text{TPR}_{g,y}^{(t)}$ and $\text{TPR}_{\sim g,y}^{(t)}$ that can result in a given TPR gender gap $\widehat{\text{Gap}}_{g,y}^{(t)}$. For example, a TPR gender gap of -0.2 might correspond to $0.6 - 0.8$ or to $0.7 - 0.9$. Moreover, different TPR pairs will result in different percentages of people with gender g in occupation y at time $t + 1$. The bands in Figure 6 therefore reflect these differences.

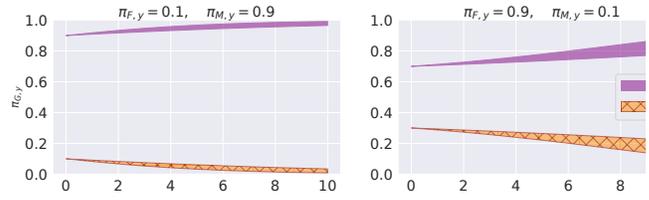


Figure 6: Simulations of compounding imbalances using the DNN representation. Each subplot corresponds to a different initial gender imbalance and shows $\pi_{g,y}^{(t)}$ for $t = 0, \dots, 10$.

william henry gates iii (born october 28 , 1955) is an american business magnate , investor , author , philanthropist , humanitarian , and principal founder of microsoft corporation . during his career at microsoft , gates held the positions of chairman , ceo and chief software architect , while also being the largest individual shareholder until may 2014 . in 1975 , gates and paul allen launched microsoft , which became the world 's largest pc software company . gates led the company as chief executive officer until stepping down in january 2000 , but he remained as chairman and created the position of chief software architect for himself . in june 2006 , gates announced that he would be transitioning from full-time work at microsoft to part-time work and full-time work at the bill & melinda gates foundation , which was established in 2000 .

Figure 7: Visualization of the DNN’s per-token attention weights. Predicted label (i.e., occupation): *software engineer*.

Attention to gender. The DNN’s per-token attention weights allow us to understand proxy behavior that occurs in the absence of explicit gender indicators. The attention weights indicate which tokens are most predictive. For example, Figure 7 depicts the per-token attention weights from the occupation classifier for the DNN representation when predicting Bill Gates’ occupation from an excerpt of his biography on Wikipedia; the larger the weight, the stronger the color. The attention weights for the words *software* and *architect* are very large, and the DNN predicts *software engineer*.

In order to understand proxy behavior that occurs in the absence of explicit gender indicators, we first used the subsampled testing split, described above, to obtain per-token attention weights from the gender classifier for the DNN representation. We then used these weights to find “proxy candidates”—i.e., the words that are most predictive of gender in the absence of explicit gender indicators. Specifically, we computed the sum of the per-token attention weights for each word type, and then selected the types with the largest sums as “proxy candidates.” Across multiple runs, we found that the words *women*, *husband*, *mother*, *woman*, and *female* (ordered by decreasing total attention) were consistently “proxy candidates.”

For each “proxy candidate,” we then used our dataset’s testing split, with and without explicit gender indicators, to create histograms of the per-token attention weights from the occupation classifier for the DNN representation. These histograms represent the extent to which that “proxy candidate” is predictive of occupation, with and without gender indicators. By comparing the histograms for each “proxy candidate,” we are able to identify words that are used as proxies for gender in the absence of explicit gender indicators: if there is a big difference between the histograms, then the “proxy candidate” is likely a proxy. Figure 8 shows per-occupation histograms for the word *women*, with (left) and without (right) explicit gender indicators. It is clear that in the absence of explicit gender indicators, the classifier has larger attention weights

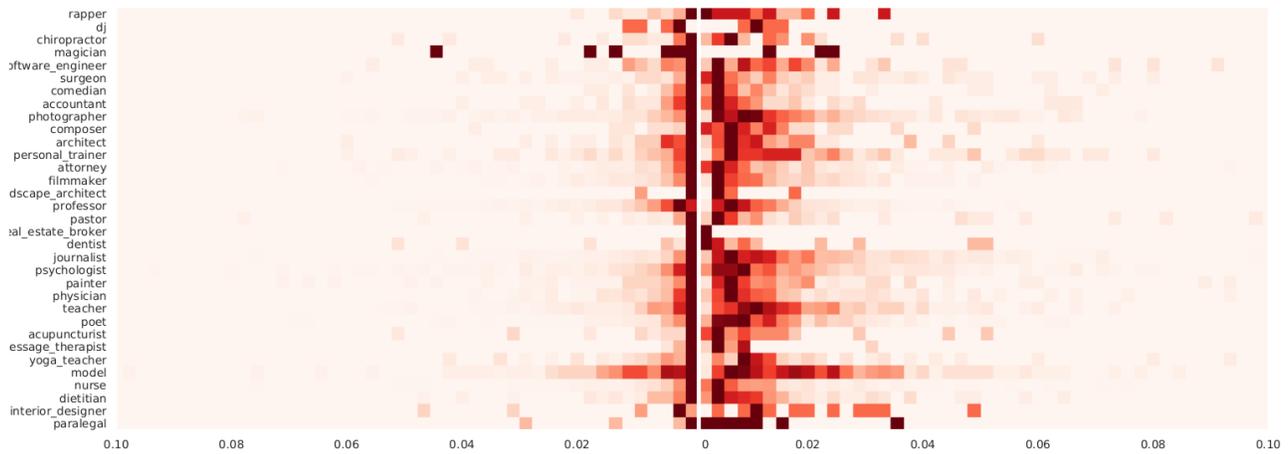


Figure 8: Per-occupation histograms of the per-token attention weights from the DNN representation’s occupation classifier for the word *women*, with (left) and without (right) explicit gender indicators; occupations are ordered by TPR gender gap.

for the word *women* for all occupations. We see similar behavior for the other “proxy candidates,” suggesting that the classifier uses proxies for gender in the absence of explicit gender indicators.

The occupations in Figure 8 are ordered by TPR gender gap from negative to positive. For occupations in the middle, where there are small or no TPR gender gaps, the classifier still has non-zero attention weights for the word *women*. This means that using gender information does not necessarily lead to a TPR gender gap. We also note that it’s possible that the classifier is using gender information to differentiate between occupations with very different gender imbalances that are otherwise similar, such as physician and surgeon.

6 DISCUSSION AND FUTURE WORK

In this paper, we presented a large-scale study of gender bias in occupation classification using a new dataset of hundreds of thousands of online biographies. We showed that there are significant TPR gender gaps when using three different semantic representations: bag-of-words, word embeddings, and deep recurrent neural networks. We also showed that the correlation between these TPR gender gaps and existing gender imbalances in occupations may compound these imbalances. By performing simulations, we demonstrated that compounding imbalances are especially problematic if people repeatedly encounter occupation classifiers because the underrepresented gender will become even further underrepresented.

Recently, Dwork and Ilvento [19] showed that fairness does not hold under composition, meaning that if two classifiers are individually fair according to some fairness metric, then the sequential use of these classifiers will not necessarily be fair according the same metric. One interpretation of our finding regarding compounding imbalances is that unfairness holds under composition. Understanding why this is the case, especially given that fairness does not hold under composition, is an interesting direction for future work.

We found that the TPR gender gaps are reduced by “scrubbing” explicit gender indicators, while the classifiers’ overall accuracies remain roughly the same. This constitutes an empirical example where there is little tradeoff between promoting fairness—in this case by “scrubbing” explicit gender indicators—and performance.

This also constitutes an empirical example where fairness is improved by “scrubbing” sensitive attributes, contrary to other examples in the literature [29]. That said, in the absence of explicit gender indicators, we did find that (1) we were able to predict the gender of a biography’s subject better than random, even when controlling for occupation; (2) significant TPR gender gaps remain for some occupations; (3) there is still a positive correlation between the TPR gender gap for an occupation and the gender imbalance in that occupation, so existing gender imbalances may be compounded. These findings indicate that there are differences between men’s and women’s online biographies other than explicit gender indicators. These differences may be due to the ways that men and women represent themselves or due to men and women having different specializations within an occupation. Our findings highlight both the risks of using machine learning in a high-stakes setting and the difficulty of trying to promote fairness by “scrubbing” sensitive attributes.

Our future work will focus primarily on understanding how best to mitigate TPR gender gaps and compounding imbalances in online recruiting and automated hiring. Finally, although we focused on gender bias, we note that other biases, such as those involving race or socioeconomic status, may also be present in occupation classification. Our methodology and analysis approach may prove useful for quantifying such biases, provided relevant group membership information is available. Moreover, quantifying such biases is an important direction for future work—it is likely that they exist and, in the absence of evidence that they do not, online recruiting and automated hiring run the risk of compounding prior injustices.

REFERENCES

- [1] Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2016. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. *arXiv preprint arXiv:1608.04207* (2016).
- [2] Kristen M Altenburger, Rajlakshmi De, Kaylyn Frazier, Nikolai Avteniev, and Jim Hamilton. 2017. Are There Gender Differences in Professional Self-Promotion? An Empirical Case Study of LinkedIn Profiles Among Recent MBA Graduates. In *ICWSM*. 460–463.
- [3] Ian Ayres. 2002. Outcome tests of racial disparities in police practices. *Justice research and Policy* 4, 1-2 (2002), 131–142.
- [4] Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *Cal. L. Rev.* 104 (2016), 671.

- [5] Marianne Bertrand and Esther Dufo. 2017. Field Experiments on Discrimination. In *Handbook of Economic Field Experiments*. Vol. 1. Elsevier, 309–393.
- [6] Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review* 94, 4 (2004), 991–1013.
- [7] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. *arXiv preprint arXiv:1805.01788* (2018).
- [8] Su Lin Blodgett and Brendan O'Connor. 2017. Racial Disparity in Natural Language Processing: A Case Study of Social Media African-American English. *arXiv preprint arXiv:1707.00061* (2017).
- [9] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [10] 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*. 4349–4357.
- [11] Victoria L Brescoll. 2011. Who takes the floor and why: Gender, power, and volubility in organizations. *Administrative Science Quarterly* 56, 4 (2011), 622–641.
- [12] Toon Calders and Indrė Žliobaitė. 2013. Why unbiased computational processes can lead to discriminative decision procedures. In *Discrimination and privacy in the information society*. Springer, 43–57.
- [13] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [14] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2018. Ranking with fairness constraints. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*.
- [15] Na Cheng, Rajarathnam Chandramouli, and KP Subbalakshmi. 2011. Author gender identification from text. *Digital Investigation* 8, 1 (2011), 78–88.
- [16] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).
- [17] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2017. Measuring and Mitigating Unintended Bias in Text Classification. (2017).
- [18] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. ACM, 214–226.
- [19] Cynthia Dwork and Christina Ilvento. 2018. Fairness Under Composition. *arXiv preprint arXiv:1806.06122* (2018).
- [20] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Mark DM Leiserson. 2018. Decoupled classifiers for group-fair and efficient machine learning. In *Conference on Fairness, Accountability and Transparency*. 119–133.
- [21] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115, 16 (2018), E3635–E3644.
- [22] Sahin Cem Geyik and Krishnamurthy Kenthapadi. October 2018. Building Representative Talent Search at LinkedIn. (October 2018). LinkedIn engineering blog post, Available at <https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin>.
- [23] Donna K Ginther and Shulamit Kahn. 2004. Women in economics: Moving up or falling off the academic career ladder? *Journal of Economic perspectives* 18, 3 (2004), 193–214.
- [24] Claudia Goldin and Cecilia Rouse. 2000. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American economic review* 90, 4 (2000), 715–741.
- [25] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*. 3315–3323.
- [26] Deborah Hellman. 2018. Indirect Discrimination and the Duty to Avoid Compounding Injustice. *Foundations of Indirect Discrimination Law, Forthcoming* (2018).
- [27] Niki Kilbertus, Mateo Rojas Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. Avoiding discrimination through causal reasoning. In *Advances in Neural Information Processing Systems*. 656–666.
- [28] Pauline T Kim. 2016. Data-driven discrimination at work. *Wm. & Mary L. Rev.* 58 (2016), 857.
- [29] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. 2018. Algorithmic fairness. In *AEA Papers and Proceedings*, Vol. 108. 22–27.
- [30] Moshe Koppel, Shlomo Argamon, and Anat Rachel Shimoni. 2002. Automatically categorizing written texts by author gender. *Literary and Linguistic Computing* 17, 4 (2002), 401–412.
- [31] Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhres, and Armand Joulin. 2018. Advances in Pre-Training Distributed Word Representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- [32] David Niven and Jeremy Zilber. 2001. Do women and men in congress cultivate different images? Evidence from congressional web sites. *Political Communication* 18, 4 (2001), 395–405.
- [33] Devin G Pope and Justin R Sydnor. 2011. Implementing anti-discrimination policies in statistical profiling models. *American Economic Journal: Economic Policy* 3, 3 (2011), 206–31.
- [34] Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. *arXiv preprint arXiv:1804.09301* (2018).
- [35] Heather Sarsons. 2015. Gender differences in recognition for group work. *Harvard University Working Paper* (2015).
- [36] Heather Sarsons. 2017. Interpreting signals in the labor market: evidence from medical referrals. *Job Market Paper* (2017).
- [37] David G Smith, Judith E Rosenstein, Margaret C Nikolov, and Darby A Chaney. 2018. The Power of Language: Gender, Status, and Agency in Performance Evaluations. *Sex Roles* (2018), 1–13.
- [38] Rachael Tatman. 2017. Gender and Dialect Bias in YouTube’s Automatic Captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*. 53–59.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- [40] Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. 22.
- [41] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1480–1489.
- [42] Meike Zehlke, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA*IR: A fair top-k ranking algorithm. In *Proceedings of the ACM Conference on Information and Knowledge Management*. 1569–1578.
- [43] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning fair representations. In *International Conference on Machine Learning*. 325–333.
- [44] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender bias in coreference resolution: Evaluation and debiasing methods. *arXiv preprint arXiv:1804.06876* (2018).