

Good Secretaries, Bad Truck Drivers? Occupational Gender Stereotypes in Sentiment Analysis

Jayadev Bhaskaran

ICME, Stanford University, USA
jayadevbhaskaran@gmail.com

Isha Bhallamudi

Dept. of Sociology, UC Irvine, USA
isha.b@uci.edu

Abstract

In this work, we investigate the presence of occupational gender stereotypes in sentiment analysis models. Such a task has implications for reducing implicit biases in these models, which are being applied to an increasingly wide variety of downstream tasks. We release a new gender-balanced dataset¹ of 800 sentences pertaining to specific professions and propose a methodology for using it as a test bench to evaluate sentiment analysis models. We evaluate the presence of occupational gender stereotypes in 3 different models using our approach, and explore their relationship with societal perceptions of occupations.

1 Motivation

Social Role Theory (Eagly and Steffen, 1984) shows that our ideas about gender are shaped by observing, over time, the roles that men and women occupy in their daily lives. These ideas can crystallize into rigid stereotypes about how men and women ought to behave, and what work they can and cannot do. Gendered stereotypes are powerful precisely for this reason: they define desirable and expected traits, roles and behaviors in people, and go beyond description to prescription. Such biases from the social world, when they map onto machine learning models, serve to reinforce and propagate stereotypes further.

In this paper, we look specifically at occupational gender stereotypes in the context of sentiment analysis. Sentiment analysis is increasingly being applied for recruitment, employee retention and job satisfaction in the corporate world (Costa and Veloso, 2015). Given the prevalence of occupational gender stereotypes, our study primarily deals with the question of whether sentiment analysis models display and propagate these stereotypes. To contextualize and ground our study, we

first provide a summary of the relevant sociological literature on occupational gender stereotypes.

1.1 Background

Sociological studies as early as 1975 (Shinar, 1975) investigate gender stereotypes of occupations, and rank occupations in terms of how “masculine”, “feminine” or neutral they are perceived to be. Cejka and Eagly (1999) successfully predicted the gender distribution of occupations based on beliefs about how specific *gender-stereotypical* attributes (such as “masculine physical”) contribute to occupational success. Such beliefs - that success in a *male* dominated profession, for example, requires *male-specific* traits - directly contribute to sex segregation in occupations. The study also found that high occupational prestige and wages are *strongly correlated* with masculine images. Together, this goes to show that occupational structure is deeply shaped by gender. More recently, Haines et al. (2016) investigate how and whether gender stereotypes have changed between 1983 and 2014, and find conclusive evidence that occupational gender stereotypes have *persisted strongly* through the ages and remain stable. There is ample sociological evidence to show that occupational gender stereotypes have *not* undergone substantial modification since the entry of women into the workplace, and that they remain pervasive and widely held by both men and women (Glick et al., 1995; Haines et al., 2016).

Since occupational gender stereotypes are shaped by subjective factors and *not* objective reality, they remain resistant to contrary evidence. Theories such as the backlash hypothesis (Rudman and Phelan, 2008) further explain their persistence: this theory shows how women in the workplace must disconfirm female stereotypes in order to be perceived as competent leaders, yet traits of ambition and capability in women evoke negative

¹Link to dataset: <https://bit.ly/2HLSKnf>

reactions which present a barrier to every level of occupational success.

The implications of occupational gender stereotypes are profound. Children and adolescents are particularly sensitive to gendered language used to describe occupations and form rigid occupational gender stereotypes based on this (Vervecken et al., 2013). In adults, occupational gender stereotypes directly contribute to the existence of unequal compensation and discriminatory hiring. They also lead to self-fulfilling prophecies: for instance, individuals may not apply to certain jobs in the first place because they think they don’t fit the gender stereotype for occupational success in that field (Kay et al., 2015).

In the following section, we discuss relevant prior work on gender bias from the NLP literature. In Section 3 we describe our methodology, dataset, and experiments in greater detail. In Section 4, we present and analyze our results, and finally, Section 5 describes possible directions of future work and concludes².

2 Prior Work

Word embeddings have been the bedrock of neural NLP models ever since the arrival of `word2vec` (Mikolov et al., 2013), and a variety of topics related to biases with word embeddings have been studied in prior literature. Garg et al. (2018) show the presence of stereotypes in word embeddings through the ages, while Bolukbasi et al. (2016) demonstrate explicit examples of social biases that are introduced into word embeddings trained on a large text corpus. Prior work has also dealt with occupational gender stereotypes in different areas of NLP. Caliskan et al. (2017) formulate a method to test biases (including gender stereotypes) in word embeddings, while Rudinger et al. (2018) investigate such stereotypes in the context of coreference resolution. There have also been efforts to *debias* word embeddings (Bolukbasi et al., 2016) and come up with *gender neutral* word embeddings (Zhao et al., 2018). These efforts, however, have attracted criticism suggesting that they do not actually *debias* embeddings but instead *re-distribute* the bias across the embedding landscape (Gonen and Goldberg, 2019).

Recent trends have been towards replacing fixed word embeddings with large pretrained *contextual*

²Source code for this paper: github.com/jayadevvhaskaran/gendered-sentiment




Figure 1: Simple diagram of our task definition.

representations as building blocks for NLP tasks. The rise of this paradigm is characterized by the use of language models for pretraining, exemplified by models such as ELMo (Peters et al., 2018), ULMFit (Howard and Ruder, 2018), GPT (Radford, 2018), and BERT (Devlin et al., 2018).

These models have shown marked improvements over word vector based approaches for a variety of tasks. However, their complexity leads to a tradeoff in terms of interpretability. Recent works have investigated gender biases in such deep contextual representations (May et al., 2019; Basta et al., 2019) as well as their applications to coreference resolution (Zhao et al., 2019; Webster et al., 2018); however, no prior work has dealt with such models in the context of occupational gender stereotypes in sentiment analysis.

Kiritchenko and Mohammad (2018) introduce the Equity Evaluation Corpus, a dataset used for measuring racial and gender biases in sentiment analysis-like systems. It was initially used to evaluate systems that predicted emotion and valence of Tweets (Mohammad et al., 2018). We use a similar approach to create a new dataset for measuring gender differences with a specific focus on occupational gender stereotypes. Our approach is *model-independent* and can be used for *any* sentiment analysis system, irrespective of model complexity.

3 Methodology

We create a dataset of 800 sentences, each with the following structure: **noun is a/an profession**. Here, *noun* corresponds to a *male* or *female* noun phrase, such as “This boy”/“This girl”, and *profession* is one of 20 different professions. Each sentence is an assertion of fact, and by itself does not seek to exhibit either positive or negative sentiment. Our dataset is balanced across genders and has 20 noun phrases for each gender, leading to a total of 400 sentences per gender.

The rationale behind our selection of the 20 professions is to include a variety of gender distribution characteristics and occupation types, in correspondence with US Current Population Survey 2018 (CPS) data (Current Population Survey, 2018) and prior literature (Haines et al., 2016). We select 5 professions that are male-dominated (*truck driver, mechanic, pilot, chef, soldier*) and 5 that are female-dominated (*teacher, flight attendant, clerk, secretary, nurse*) - with domination meaning greater than 70% share in the job distribution. Next, we add professions that are slightly male-dominated (*scientist, lawyer, doctor*) and slightly female-dominated (*writer, dancer*), with slight domination meaning a 60 – 65% share in the job distribution. We also add *professor*, which does not have a clear definition as per CPS but has been known to have different gender splits at senior and junior levels. Finally, we include two professions that show an approximately neutral divide (*tailor, gym trainer*) and two which have experienced significant changes in their gender distribution over time (*baker, bartender*), with an increasing female representation in recent times (Haines et al., 2016). As mentioned previously, we also select our set of occupations with an eye towards representing a range of occupation types.

We evaluate 3 sentiment analysis models through our experiments. Each model is trained on the Stanford Sentiment Treebank 2 `train` dataset (Socher et al., 2013), which contains phrases from movie reviews along with binary (0/1) sentiment labels. We then evaluate each model on our new corpus and measure the difference in mean predicted positive class probabilities between sentences with male nouns and those with female nouns. We test 3 hypotheses (one for each model), with the *null hypotheses* indicating *no difference* in means between sentences with *male* and *female* nouns. Fig. 1 illustrates our experimental setup.

Our evaluation methodology is very similar to that used in Kiritchenko and Mohammad (2018). For each system, we predict the positive class probability for each sentence. We then apply a *paired* t-test (since each pair contains a *male* and *female* version of the same template sentence) to measure if the mean predicted positive class probabilities are different across genders, using a significance level of 0.01. Since we test *three* hypotheses (one for each system), we apply Bonferroni correction (Bonferroni, 1936) to the *p*-values

that we obtain. In other words, the null hypothesis is rejected only for calculated *p*-values less than 0.01/3. We note that we do not perform any correction to account for the fact that the sentences within each gender are *not* iid, and only vary in the `noun` and `profession` words.

The 3 models that we evaluate are as follows:

- **M.1:** Bag-of-words + Logistic Regression (baseline): We build a simple bag-of-words model, apply tf-idf weighting, and use logistic regression (implemented using `scikit-learn` (Pedregosa et al., 2011)) to classify sentiment. This model is a very simple approach that has nevertheless been found to work well in practice for sentiment analysis tasks, and we use it as our baseline model.
- **M.2:** BiLSTM: We use a bidirectional LSTM implemented in `Keras` (Chollet et al., 2015) to predict sentiment. The words in a sentence are represented by 300-dimensional GloVe embeddings (Pennington et al., 2014). This model is more sophisticated than the baseline and captures some contextual information and long-term dependencies (Hochreiter and Schmidhuber, 1997). This model also allows us to investigate gender differences that might be introduced through word embeddings, as described in Bolukbasi et al. (2016).
- **M.3:** BERT (Devlin et al., 2018): We use a pretrained (uncased) BERT-Base model³ and finetune it the SST-2 dataset. This shows near state-of-the-art performance on a wide variety of NLP tasks, including sentiment analysis (Devlin et al., 2018).

While analysing the results of our experiments, we measure overall predicted mean positive probabilities (across genders) for each of the 20 professions in our newly created dataset, to identify which professions are rated as *high-sentiment* by these models. This helps us investigate relationships between societal perceptions of occupations and corresponding sentiment predictions from the models.

We also examine differences in sentiment among equivalent gender pairs (such as *bachelor* and *spinster*) for the 20 pairs in our dataset, to investigate differences in predicted sentiment between different sets of male/female noun pairs.

³Model source: <https://bit.ly/2S8w6Jt>

Model	Dev Acc.	F - M
M.1 (BoW+LogReg)	0.827	0.035**
M.2 (BiLSTM)	0.841	0.077**
M.3 (BERT)	0.930	-0.040**

Table 1: Results. Dev Acc. represents accuracy on SST-2 dev set. **F - M** represents difference between means of predicted positive class probabilities for sentences with *female* nouns and sentences with *male* nouns. ** denotes statistical significance with $p < 0.01$ (after applying Bonferroni correction).

Finally, we examine differences between *male* and *female* nouns for each individual occupation, to understand which occupations are susceptible to gender stereotyping.

4 Results/Analysis

The main results of our experiments are shown in Table 1. Our null hypothesis is that the predicted positive probabilities for *female* and *male* sentences have identical means. We notice that **M.1** (Bag-of-words + Logistic Regression) and **M.2** (BiLSTM) show a statistically significant difference between the two genders, with higher predicted positive class probabilities for sentences with *female* nouns. This effectively represents the biases seen in the SST-2 train dataset. The dataset has 1182 sentences containing a *male* noun with a mean sentiment of **0.535**, and 601 sentences containing a *female* noun with a mean sentiment of **0.599**. Thus, biases present in training data can get propagated through machine learning models, and our approach can help identify these.

On the contrary, **M.3** (BERT) shows that sentences with *male* nouns have a statistically significant higher predicted positive class probability than sentences with *female* nouns. One possible reason for this might be biases that propagate from the pretraining phase in BERT. This finding indicates a promising direction of future work: investigating the effects of gender biases in the large pretraining corpus versus those in the smaller fine-tuning corpus (in our case, the SST-2 train dataset).

4.1 Social Stereotypes of Occupations

We now look at mean distributions of positive class probability (across genders) for each profession, as shown in Table 2. We notice that *secretary* shows up as a high positive sentiment profes-

Model	Top 3 professions
BoW+LogReg	Secretary, Teacher, Writer
BiLSTM	Dancer, Secretary, Scientist
BERT	Scientist, Chef, Dancer
Model	Bottom 3 professions
BoW+LogReg	Truck Dr., Fl. Att., (many)
BiLSTM	Truck Dr., Gym Tr., Nurse
BERT	Truck Dr., Clerk, Tailor

Table 2: Top 3 and bottom 3 professions per model, based on predicted positive class probability (agnostic of gender). Note: For BoW+LogReg, (many) denotes all the professions that did not appear in the SST-2 train dataset.

sion in both **M.1** and **M.2**. On further investigation, we notice that this artefact arises because of the 2002 movie *Secretary*, starring Maggie Gyllenhall, that has a number of positive reviews that form a part of the SST-2 train dataset. However, **M.3** (BERT) seems to be impervious to this, indicating that extensive pretraining could have the potential to remove certain corpus-specific effects that might have lingered in shallower models.

The profession with the lowest average sentiment score across all 3 models is *truck driver*; other low scoring professions include *clerk*, *gym trainer* and *flight attendant*. We also note that the highest scoring profession (average sentiment 0.99) with **M.3** (BERT) is *scientist* and the lowest (average sentiment 0.34) is *truck driver*, disturbingly reflective of societal stereotypes about white-collar and blue-collar jobs.

To explore this further, we look at data from the Current Population Survey of the US Bureau of Labor Statistics (Current Population Survey, 2018). Fig. 2 shows the relationship between median weekly earnings (for occupations where data is available) and average positive sentiment predicted by BERT. While there are some outliers, the figure shows a positive correlation between earnings and sentiment, indicating that the model may have incorporated societal perceptions around different occupations. We note that this is only a rough analysis, as not all occupations directly correspond to entries from the survey data.

4.2 Gendered Stereotypes

We attempt to analyze differences in gender within occupations by studying the predictions of **M.3** (BERT), which incorporates the largest amount




Figure 2: Median weekly earnings (Current Population Survey, 2018) vs. mean predicted positive probability using M.3 (BERT), per profession.

of external data. First, we analyze differences in mean positive class probability between sentences with male and female nouns for each profession. We notice that *pilot* has the highest *positive* difference between female and male noun sentences (i.e., *female is higher*), while *flight attendant* has the *most negative* difference (i.e., *male is higher*). This provides an interesting dichotomy: *pilot* is a male-dominated profession, while *flight attendant* is a female-dominated one.

To test whether these are just artefacts of generic gender bias in the model or specific to occupational gendered stereotypes, we replace *profession* with “person” to create 20 sentence pairs such as “This man/this woman is a person.”, and predict the sentiment for these 20 pairs. We notice that the difference between *female* and *male* noun sentences for the control experiment is **0.039**, showing that sentences with *female* nouns in the control group exhibit *higher* positive sentiment than those with *male* nouns. The three occupations with the most negative difference (i.e., *female* sentences have lower positive sentiment) are *flight attendant* (**-0.132**), *bartender* (**-0.126**), and *clerk* (**-0.116**). Of these, *flight attendant* (72%) and *clerk* (86%) are female-dominated professions (Current Population Survey, 2018), while *bartender* (55%) is a profession that has been shifting from male to female-dominated in recent times (Haines et al., 2016).

Finally, we study differences between corresponding pairs of female and male nouns, using predictions from M.3 (BERT). Out of the 20 pairs in our dataset, the pair with the greatest difference in mean positive class probability is *spinster* and *bachelor*, with

$\text{spinster} - \text{bachelor} = -0.404$ ($p < 0.01$). This reflects societal perceptions of *spinster* as someone who is characterized as alone, lonely and resembling an “old maid”, versus *bachelor* as someone who might be young, carefree and fun-loving (Nieuwets, 2015). This is an example of *semantic pejoration* seen in society, where the female form of the noun (i.e., *spinster*) gradually acquires a negative connotation. Notably, this pejorative behavior may have also leaked into the model, reflecting societal gender stereotypes.

5 Conclusion/Future Work

In this paper, we introduce a new dataset that can be used to test the presence of occupational gender stereotypes in *any* sentiment analysis model. We then train 3 sentiment analysis models and evaluate them using our dataset. Following that, we analyze our results, exploring social stereotypes of occupations as well as gendered stereotypes. We find that all 3 models that we study exhibit differences in mean predicted positive class probability between genders, though the directions vary. We also see that simpler models may be more susceptible to biases seen in the training dataset, while deep contextual models may exhibit biases potentially introduced during pretraining.

One promising avenue for future work is to explore occupational stereotypes in deep contextual models by analyzing their training corpora. This could also help identify techniques to mitigate biases in such models, since they could be relatively impervious to biases introduced by fine-tuning (especially on smaller datasets).

From a sociological perspective, we plan to investigate occupational gender stereotypes in downstream applications such as automated resume screening. Such a task assumes greater importance with the increased use of these systems in today’s world. There is prior work on ethnic bias in such tools (Derous and Ryan, 2018), and we believe that there is significant value in exploring and characterizing gender biases in these systems.

Acknowledgements

We thank Johan Ugander for helping motivate the initial phases of this work. We also thank the anonymous reviewers for their thoughtful feedback and suggestions.

References

- Christine Basta, Marta R. Costa-jussà, and Noe Casas. 2019. [Evaluating the Underlying Gender Bias in Contextualized Word Embeddings](#). *arXiv e-prints*.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. [Man is to Computer Programmer As Woman is to Home-maker? Debiasing Word Embeddings](#). In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 4356–4364, USA. Curran Associates Inc.
- Carlo Bonferroni. 1936. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Mary Ann Cejka and Alice H. Eagly. 1999. [Gender-Stereotypic Images of Occupations Correspond to the Sex Segregation of Employment](#). *Personality and Social Psychology Bulletin*, 25(4):413–423.
- François Chollet et al. 2015. Keras. <https://keras.io>.
- Andre Costa and Adriano Veloso. 2015. [Employee analytics through sentiment analysis](#). In *Brazilian Symposium on Databases*.
- Current Population Survey. 2018. [39. Median weekly earnings of full-time wage and salary workers by detailed occupation and sex](#). Bureau of Labor Statistics, United States Department of Labor.
- Eva Derous and Ann Marie Ryan. 2018. [When your resume is \(not\) turning you down: Modelling ethnic bias in resume screening](#). *Human Resource Management Journal*, 29(2):113–130.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of Deep Bidirectional Transformers for Language understanding](#). *CoRR*, abs/1810.04805.
- Alice H Eagly and Valerie J Steffen. 1984. Gender stereotypes stem from the distribution of women and men into social roles. *Journal of personality and social psychology*, 46(4):735.
- 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):E3635–E3644.
- Peter Glick, Korin Wilk, and Michele Perreault. 1995. [Images of occupations: Components of gender and status in occupational stereotypes](#). *Sex Roles*, 32(9):565–582.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a Pig: Debiasing Methods Cover up Systematic Gender Biases in Word Embeddings But do not Remove Them](#). *arXiv e-prints*.
- Elizabeth L. Haines, Kay Deaux, and Nicole Lofaro. 2016. [The Times They Are a-Changing... or Are They Not? A Comparison of Gender Stereotypes, 1983-2014](#). *Psychology of Women Quarterly*, 40(3):353–363.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Comput.*, 9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. [Fine-tuned Language Models for Text Classification](#). *CoRR*, abs/1801.06146.
- Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. [Unequal representation and gender stereotypes in image search results for occupations](#). In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pages 3819–3828, New York, NY, USA. ACM.
- Svetlana Kiritchenko and Saif M. Mohammad. 2018. [Examining gender and race bias in two hundred sentiment analysis systems](#). *CoRR*, abs/1805.04508.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. [On measuring social biases in sentence encoders](#).
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed Representations of Words and Phrases and their Compositionality](#). In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. [Semeval-2018 Task 1: Affect in Tweets](#). In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.
- Astrid Nieuwets. 2015. *Fallen Females: On the Semantic Pejoration of Mistress and Spinster*. Bachelor’s thesis, Utrecht University.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global Vectors for Word Representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). *CoRR*, abs/1802.05365.
- Alec Radford. 2018. [Improving Language Understanding by Generative Pre-Training](#).
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*.
- Laurie A. Rudman and Julie E. Phelan. 2008. [Backlash effects for disconfirming gender stereotypes in organizations](#). *Research in Organizational Behavior*, 28:61 – 79.
- Eva H Shinar. 1975. [Sexual stereotypes of occupations](#). *Journal of Vocational Behavior*, 7(1):99 – 111.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1631–1642.
- Dries Vervecken, Bettina Hannover, and Ilka Wolter. 2013. [Changing \(S\)expectations: How gender fair job descriptions impact children’s perceptions and interest regarding traditionally male occupations](#). *Journal of Vocational Behavior*, 82(3):208 – 220.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. [Mind the GAP: A Balanced Corpus of Gendered Ambiguous Pronouns](#). *Transactions of the Association for Computational Linguistics*, 6:605617.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. [Gender Bias in Contextualized Word Embeddings](#). *CoRR*, abs/1904.03310.
- Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. [Learning Gender-Neutral Word Embeddings](#). *CoRR*, abs/1809.01496.