

Professional Presentation and Projected Power: A Case Study of Implicit Gender Information in English CVs

Jinrui Yang^{1*}, Sheilla Njoto^{2*♥}, Marc Cheong¹, Leah Ruppner², and Lea Frermann¹

¹School of Computing and Information Systems, University of Melbourne

²School of Social and Political Sciences, University of Melbourne

{jinruiy, snjoto}@student.unimelb.edu.au

{marc.cheong, leah.ruppner, lea.frermann}@unimelb.edu.au

Abstract

Gender discrimination in hiring is a pertinent and persistent bias in society, and a common motivating example for exploring bias in NLP. However, the manifestation of gendered language in application materials has received limited attention. This paper investigates the framing of skills and background in CVs of self-identified men and women. We introduce a data set of 1.8K authentic, English-language, CVs from the US, covering 16 occupations, allowing us to partially control for the confound occupation-specific gender base rates. We find that (1) women use more verbs evoking impressions of low power; and (2) classifiers capture gender signal even after data balancing and removal of pronouns and named entities, and this holds for both transformer-based and linear classifiers.

1 Introduction

In this paper, we study word choice and implied power and agency in curriculum vitae (CVs) authored by men and women, combining lines of research that emerged from a long research tradition in both the social sciences and, more recently, natural language processing (Carli, 1990; Lakoff, 1975; Glick and Fiske, 2018). From a sociology perspective, it has been suggested that choices of words are influenced by the social status of the respective genders at a given moment in society (Talbot, 2019). Women are known to use more communal forms of words and emotional connotations than men (Brownlow et al., 2003; Leaper and Ayres, 2007; Newman et al., 2008), and that such choices reflect the different levels of power and influence both politically and economically (Talbot, 2019; Leaper and Ayres, 2007). Conversely, the choice of language impacts how the reader *perceives* the entity described in the text. In particular, the choice of

verbs has been suggested as an indicator of the perceived levels of *power* and *agency* of the described entity (Sap et al., 2017).

Organisational scholars have long documented gender discrimination in employment (Booth and Leigh, 2010; Heilman, 2012; Steinpreis et al., 1999). Sociological studies have repeatedly shown that women are evaluated more harshly than men especially in recruitment (Moss-Racusin et al., 2012; Neumark, 2010; Riach and Rich, 2006). Men tend to be assessed for their competence, while women are assessed based on characteristics (‘likeability’), even when they demonstrate the same levels of qualifications, experience and education (Rudman, 1998; Phelan et al., 2008). Gaucher et al. (2011) studied the impact of “gendered wording” in job advertisements on gender inequality in traditionally male-dominated occupations via content analysis, while De-Arteaga et al. (2019) showed how gender signal in online biographies lead to disparate performance in the task of occupation classification. Experience has shown that leaving hiring decisions to supposedly objective algorithms did not remove bias from the process – both in real-world applications like Amazon’s gender-biased automatic hiring tool (Bogen, 2019), as well as a surge in research on showing and alleviating bias in NLP models (Sun et al., 2019).

We present a data set of 1.8K human-written CVs and study differences in word choice and framing between men and women, and the extent to which classifiers are susceptible to gendered language. Unlike prior studies which were either occupation-specific (Parasurama and Sedoc, 2022) or used proxy data like online biographies (De-Arteaga et al., 2019), we inspect application materials directly and cover 16 occupations (Appendix B) which allows us to study gender differences while partially controlling for the confound of occupation-specific base rates. However, we find that even within occupations, confounds remain as

* Equal contribution

♥ Corresponding author

women tend to occupy lower ranking positions and men and women cluster in different types of fine-tuned jobs within an occupation category (Section 3).

Our CV authors provide self-identified gender as part of our screening questions.¹ Due to the very low number of "Other/non-binary" responses (0.01%), we here only consider people self-identifying as male or female. We acknowledge that treating gender as a binary phenomenon is an oversimplification (Guo and Caliskan, 2020), but stress that our methodology extends to more inclusive notions of gender, and hope that our study inspires future work in this direction.

After presenting our data set (Section 2), we investigate gender signals in CVs in terms of overall word choice (Section 3); implied associations with power and agency (Section 4); and predictive models' sensitivity to gender when trained on data from which gender-indicative signals were removed to different extents (Section 5).

2 Dataset collection

On Prolific², we hired 2,000 participants (50% women) who were (1) US American and live and work in the US; and (2) in full-time employment. After answering a number of screening questions, participants composed a CV "pretending that you were applying for your next promotion". We specifically asked our participants to copy from their existing CV, instead of write an entirely new CV to mimic real-world CVs as closely as possible. It was encouraged to anonymize information wherever possible, but otherwise craft a CV as realistic as possible given their current situation. For a uniformed structure, we segmented the CV submission into five parts, each as a free-text box: (1) an optional professional summary/career objective, (2) professional experience, (3) education; (4) skills and attributes; (5) optional certifications/qualifications.

Quality control and preprocessing We removed responses based on very short (long) response times and non-English text (~ 10%), retaining 1,789 CVs (50.5% female). We tokenized, lemmatized and POS-tagged all text, removed stop words, and concatenated the five CV segments. We identified pro-

nouns and named entities.³ All preprocessing was done using SpaCy's default English models.

Data sharing In line with our IRB approval, we release a deidentified version of our data set to individual researchers. Further details are in the Ethics Statement. Appendix A contains the consent form.

3 Gender-associated word choice

We qualitatively analyze gender-associated word choice in 6 (out of 16) occupations: 2× female dominated (Education, Healthcare); 2× male dominated (Computer/math, Management) and 2× balanced (Business/finance, Sales).

We first obtain the top 1% of TFIDF-ranked unigrams for both men's (M) and women's (F) CVs. We then retain terms in these two sets unique to M (and conversely, unique to F) as terms highly associated with only one gender. Due to space constraints, we present the full results in Table 4 in Appendix C.

In men-dominated occupations, men-associated terms are 'scientific' (engineer, developer, database), or relate to leadership/tactics (leadership, planning); women-associated terms relate to interpersonal skills (community, communication, social). For women-dominated occupations, terms more likely to be used by women include those related to support and teamwork (help, assistant, aid); whereas men use terms which are again 'scientific' and exhibiting leadership (physician, lead, manager).

The overall, across-occupation, pattern is not dissimilar to the occupation-stratified analyses above. This is consistent with sociological studies which have shown that men are often assessed by their competence and leadership qualities, whereas women are often assessed by their 'likeability' (i.e., their personal characters) (Eagly and Karau, 2002). On the contrary, women who show ambition and competitiveness are often penalised for violating traditional feminine stereotypes (Phelan et al., 2008). Such biased judgments are likely to discourage women to use words to describe their expertise and use more communal words instead.

Note, however, that these differences arise not only from lexical choice, but also from real world differences *within* an occupational group,

¹Participants chose from: [Man, Woman, Other/non-binary, prefer not to say].

²<https://www.prolific.co/>

³Including all entity types covered by SpaCy's default entity tagger.

where the genders distribute differently across work tasks and finer-grained roles: women tend to have lower-ranking jobs, and specific occupations within the broader groups will exhibit different gender skews. Results for the Education occupation illustrate this well, where men-associated terms are dominated by technology and leadership (microsoft, lead, technology), while women-associated terms focus on early education and support (child, elementary, social). See underlined terms in Table 4 for further examples.

4 Power and Agency in CVs

Do men and women differ in the way they present themselves in a CV? We compare the extent of *power* and *agency* implied in the verbs used by male and female applicants. We apply Sap et al. (2017)’s connotation frames of power and agency, which associate verbs with the reactions they evoke the reader (Rashkin et al., 2016). By focussing on verbs, we abstract away from (named) entities with a strong occupation association and focus on self-presentation (Goffman, 1959). We consider all transitive verbs in CVs. Given the content (primarily focused on the author) and style (listings, incomplete sentences) of CVs, we assume that the agent of every verb is the author. The **power** dimension distinguishes verbs where the agent (subject) has more ($A>T$; ‘lead’), less ($A<T$; ‘assist’), or equal ($A=T$; ‘care’) power to the theme (object). The **agency** dimension categorizes verbs as high (+; ‘support’), low (-; ‘wait’) or neutral (*neu*; ‘access’) agency. We use Sap et al. (2017)’s frame-labeled data set of 2K English verbs. 48% of verb types in our CVs are in the labeled data set (conversely, 57% of frame-labeled verbs occur in our CVs). The numbers are comparable across genders.

Overall label distribution We restrict our analysis to CVs with at least 10 and at most 100 verbs ($N=1503$, 53% women) to reduce the impact of outliers,⁴ and retrieve the power and agency label of each verb that is included in the labeled set. Figure 1 shows the overall distribution of power and agency levels in our data set. Consistent with prior work (Sap et al., 2017), and unsurprising given the data domain, we observe a dominance of agent-power and high agency verbs.

Gender differences We measure the statistical dependence of power/agency levels (independent

⁴Noting that the results hold with all data points included.

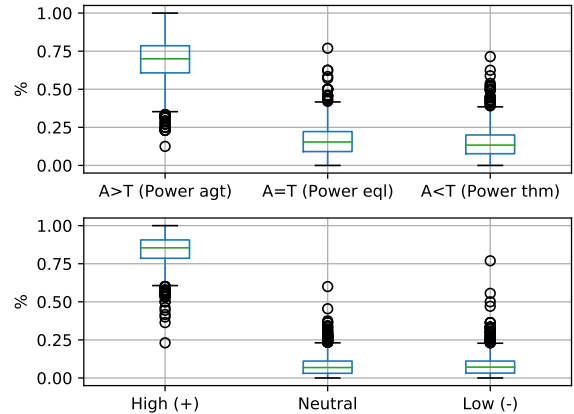


Figure 1: Distribution of power (top) and agency (bottom) verbs in our data. Each dot corresponds to one CV.

variables) on the gender of the CV author (dependent variable) fitting a logistic regression.⁵ Each CV is represented as a count vector over the 3 power and 3 agency categories, while controlling for CV length (number of words) and occupation (cf. Appendix B). We standardized features for better interpretability of β , and coded Man as 0, and Woman as 1. Table 1 shows that women use equal power ($A=T$) and theme power ($A<T$) verbs significantly more often than men. Examples for equal power verbs that are more frequently used by females than males are {complete, perform, analyze, assess}, and for theme power verbs {assist, learn, need, serve}. Gender difference for use of agent-power verbs is insignificant.

We also find that both high (+) and low (-) agency verbs are more frequently used by men.⁶ Male-associated positive agency verbs include {employ, reduce, rate, acquire, exceed} while male-associated negative agency verbs include {address, expect, stay, relate}. There are no significant differences in the use of neutral agency verbs.

4.1 Discussion

Women use more verbs in CVs that associate low power with the agent (CV author). Broadly, this agrees with prior work revealing that women are portrayed as less powerful in fiction movies (Sap et al., 2017). It also links into results from sociology revealing that leadership qualities (strength,

⁵https://www.statsmodels.org/dev/generated/statsmodels.discrete.discrete_model.Logit.html

⁶Noting that the significance for the differences in agency do not hold after Holm–Bonferroni correction (Holm, 1979) for multiple comparisons.

Feature	β	p	
Power($A > T$)	4.57	0.140	
Power($A = T$)	10.78	0.006 **	F
Power($A < T$)	15.25	0.000 **	F
Agency(+)	-6.35	0.032 *	M
Agency(<i>neu</i>)	2.34	0.589	
Agency(-)	-11.72	0.007 *	M

Table 1: Association of power/agency with binary CV author gender via coefficients (β) and significance estimates (p) of a logistic regression, after controlling for CV length and occupation. The final column indicates direction of association (male=0, female=1). * indicates statistical significance at $p < 0.05$, while ** additionally confirms significance after Holm–Bonferroni correction for multiple comparisons.

assertiveness) are evaluated more positively in men, than in women (Eagly and Karau, 2002; Rudman and Glick, 2001). Similarly, men are often rated more favorable than women given the same qualification, which might lead women to elaborate more on their education and qualifications (Njoto et al., 2022). Indeed, we find the Education section in female-authored CVs to be on average 15% longer than in male (247 vs 214 words), the Qualification and Training section 24% longer (144 vs 116 words), while the Professional Experience sections are of similar length (+2%; 1109 vs 1034 words).

Notably, this may be because, on average, women in the US have attained a higher level of education than men (Parker, 2021). However, it also reflects that women tend to get more education for the same job, and tend to be overly qualified for similar positions. Women access more education but also need more education and training for the same job (Campbell and Hahl, 2022).

5 Gender prediction from CV text

Sections 3 and 4 explored gender-specific content framing differences CVs which may impact judgment the reader (or hirer). We next quantify the susceptibility of representative predictive models to gender information in CVs. We use the task of binary author gender prediction based on the text of the CV as a diagnostic tool to assess the extent to which models can infer gender information from CVs. We explicitly caution against using this task as a ML benchmark (cf., Ethics Statement).

We test the following binary classifiers: a linear SVC with L1 regularization, which by design learns sparse and interpretable features; a lo-

Model	D-Full	D-Balanced
Random U	0.50 (± 0.00)	0.48 (± 0.00)
Majority	0.34 (± 0.00)	0.33 (± 0.00)
SVC Full	0.69 (± 0.03)	0.64 (± 0.03)
LR Full	0.72 (± 0.03)	0.66 (± 0.02)
RoBERTa Full	0.75 (± 0.02)	0.71 (± 0.03)
SVC -PER	0.69 (± 0.01)	0.61 (± 0.03)
LR -PER	0.73 (± 0.03)	0.66 (± 0.02)
RoBERTa -PER	0.75 (± 0.01)	0.57 (± 0.20)
SVC -NE	0.67 (± 0.02)	0.62 (± 0.01)
LR -NE	0.71 (± 0.02)	0.66 (± 0.02)
RoBERTa -NE	0.73 (± 0.02)	0.66 (± 0.01)

Table 2: Predicting the gender (M,F) of an author of a CV. Macro-averaged F1 score (\pm standard deviation) from 5-fold cross-validation.

gistic regression classifier (LR); and a fine-tuned RoBERTa-based classifier built on pre-trained RoBERTa uncased (Liu et al., 2019), fine-tuned for two epochs with a learning rate of 4×10^{-5} .⁷ We use TFIDF features ($|X|=5000$) for LR and SVM and plain text for RoBERTa. We include a random uniform, and a majority baseline and run all models with 5-fold cross-validation.

We test our models on three versions of our CVs.⁸ All versions are lemmatized and stopwords were removed. (1) the full data set with all lemmas from all CV sections (**Full**); (2) mask names and pronouns to remove explicit gender indicators (**-PER**); (3) remove *all* named entities (**-NE**), to abstract away from institutional information such as single-gender schools which may carry implicit information about the gender of the applicant.

Training these models on the full data sets (D-Full, N=1503) set will inevitably add a confounding factor of occupation-specific terms: most occupations are substantially gender-skewed in their workforce. To remove this confound, we create a version of each data set with a gender-balanced set of CVs *for each occupation* (D-Balanced, N=1118). We report results as macro averaged F1 scores, as presented in Table 2.

Results We test whether gendered information is encoded in classifiers trained on data with varying amounts of gender-indicative information. A perfectly gender-agnostic model would perform on par with the baselines.

Table 2 shows that all classifiers outperform the

⁷BERT uncased performed slightly below RoBERTa.

⁸Like in Section 4 we remove CVs with fewer than 10 or more than 100 verbs for consistency.

baselines substantially, both when trained on D-Full as well as on D-Balanced, where the occupation proxy gender is reduced. In line with prior work (De-Arteaga et al., 2019), we find that ‘scrubbing’ names and pronouns as explicit gender indicators (-PER) has negligible impact. Removing all named entities reduces classifier performance, but it remains well above random. RoBERTa outperforms the linear models in the Full data condition (left column), but shows unstable performance in the Balanced condition (right column) presumably due to the smaller data set leading to overfitting (note the high std in the -PER condition). Overall, the findings highlight the importance of considering gendered language signals beyond explicit indicators, i.e., that simple methods like removal of names (Manikandan, 2020) does not imply absence of gender information. Table 5 (Appendix D) lists the 20 most predictive features for the linear SVC trained on the D-Balanced, when trained on the full data (top) and the entity-redacted data (bottom). The features from the full data include entities like state names (Indianapolis, Colorado). Even after gender balancing per occupation, stereotypically associated features with women (child, and “soft” attributes like attitude, assist, document) and men (supervise, technology) emerge.

6 Discussion

We presented a data set of 1.8K authentic, US-English CVs across 16 occupations, aligned with self-reported binary gender of the author. This data set allowed us to inspect features of men- and women-specific language in CVs, while controlling for the confounding factor of occupation: most occupations are heavily gender-skewed.

This paper connects the concept of framing, i.e., influencing readers of a document through careful choice of words (Entman, 2007), with existing power discrepancies between men and women in western society in general, and the job market specifically (Rudman, 1998). We showed that women use verbs that imply lower power significantly more often than men, even after controlling for occupation. Subtle changes in word choice have been shown to impact human perception, reaction and choice (Kahneman and Tversky, 2013). In the context of *human* hiring, this suggests that (a) removing explicit gender indicators is insufficient; and (b) further support for sensitizing both hirers and applicants to subconscious bias.

We further trained classifiers to predict binary author gender based on CV text, in scenarios where gender proxy information was removed by gender-balancing the training data and/or removing named entities. We show that classifiers perform significantly above chance across all settings, confirming that subtle gender signal remains. This result is expected, and in line with prior research (De-Arteaga et al., 2019), but for the first time shown directly on data more akin to application materials presented to human and automatic hirers.

Our experiments retain a confounding factor of job type within an occupational group: within an occupation, women tend to have lower-ranking jobs; and within our 16 broad occupational groups, different specific occupations will exhibit different gender skews. In Section 3, we inspected gender-associated word choice in 6 most frequent occupations in our data set, finding that across occupations, ‘scientific’ terms (engineer, developer, database) and leadership terms (leadership, administration, planning) are more associated with male CVs; while women are more likely to mention interpersonal skills, support, or teamwork (community, communication, social, help, assistant, aid), typically associated with administrative roles. Consequently, gender signals in CVs not only originate from lexical choice, but also also reflect real-world differences in work tasks and position levels. Disentangling these factors is an important direction for future work.

In sum, we maintain that perpetuated gendered patterns embedded in CVs can bias both human and automated hiring, and that the naive use of ML methods bears the risk of exacerbating bias: by picking up spurious associations on different levels from explicit gender information (names, hobbies) to subtler word choice (the level ‘power’ or ‘agency’). Suggestions for further work include usability studies and social-psychological interventions for users of recruitment software, for *both* job applicants and decision makers. Interventions could include ‘nudges’ in the user experience flow, informing users about potential gender signals being encoded in their data, and suggestions of strategies to mitigate or minimise this. As our findings suggest, scrubbing names and entities off the CVs is not effective in de-gendering CVs for fairer recruitment decisions, and should not be used as the be-all-and-end-all in promoting fair hiring, as often is the wont of current initiatives.

Ethics Statement

This study was approved by the University of Melbourne ethics board (Human Ethics Committee LNR 3A), Reference Number 2022-22062-32741-5, and data acquisition and analysis has been taken out to the according ethical standards. Our data was collected via Prolific. The crowdworkers (annotators) in the project were paid £3.75 for a median of 11 minutes of work, which is slightly above minimum hourly wage and reflects adequate compensation for the time spent. Appendix A contains the consent form presented to annotators before the task. Prolific allows us to record information anonymously without personally identifiable data. As part of CV generation, our crowdworkers were instructed to exclude their names and the names of their affiliated organisations from their drafted CVs.

To enable future research in this area, we plan to release an anonymized and deidentified version of our data to individual researchers where names, emails, addresses, phone numbers and all named entities are redacted (the [-NE] version used in this paper). The data will contain the redacted CV text and self-identified gender label only. Interested researchers will sign an agreement form stating that they (1) will not share the data with anyone else; (2) will delete the data upon completion of the research or after 1 year whichever comes first.

This paper investigated the language differences between men and women authored CVs. Gender information was identified by the CV authors and no gender-inference was applied anywhere in the paper. We acknowledge that a binary notion of gender is not representative of the concept. In addition, we acknowledge that our study excludes a large portion of the population which does not identify to a cis-normative group. We emphasized throughout the paper that our findings hold for self-identifying men and women only, and that our methodology in principle extends to a more inclusive set of gender groups, conditioned on the availability of reliable data.

We used the task of gender prediction from CV data as a benchmark to assess the amount of gendered information retained in ML models after various strategies to remove gender proxy information. We do not endorse this task in general, and accordingly do not release pre-trained models to the public.

Acknowledgments

We thank the reviewers for their very insightful feedback. This work was partially funded by the Seed Funding scheme of the Melbourne Center for Data Science.

References

- Miranda Bogen. 2019. All the ways hiring algorithms can introduce bias. *Harvard Business Review*, 6:2019.
- Alison Booth and Andrew Leigh. 2010. Do employers discriminate by gender? a field experiment in female-dominated occupations. *Economics Letters*, 107(2):236–238.
- Sheila Brownlow, Julie A Rosamond, and Jennifer A Parker. 2003. Gender-linked linguistic behavior in television interviews. *Sex Roles*, 49(3):121–132.
- Elizabeth Lauren Campbell and Oliver Hahl. 2022. He’s overqualified, she’s highly committed: Qualification signals and gendered assumptions about job candidate commitment. *Organization Science*.
- Linda L Carli. 1990. Gender, language, and influence. *Journal of personality and social psychology*, 59(5):941.
- Maria De-Arteaga, Alexey Romanov, Hanna Wal-lach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128.
- Alice H Eagly and Steven J Karau. 2002. [Role congruity theory of prejudice toward female leaders](#). *Psychol. Rev.*, 109(3):573–598.
- Robert M Entman. 2007. Framing bias: Media in the distribution of power. *Journal of communication*, 57(1):163–173.
- Danielle Gaucher, Justin Friesen, and Aaron C Kay. 2011. [Evidence that gendered wording in job advertisements exists and sustains gender inequality](#). *J. Pers. Soc. Psychol.*, 101(1):109–128.
- Peter Glick and Susan T Fiske. 2018. The ambivalent sexism inventory: Differentiating hostile and benevolent sexism. In *Social cognition*, pages 116–160. Routledge.
- Erving Goffman. 1959. *The presentation of self in everyday life*. University of Edinburgh Social Sciences Research Centre.
- Wei Guo and Aylin Caliskan. 2020. [Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases](#).

- ME Heilman. 2012. Gender stereotypes and workplace bias. *research in organisational behaviour*, 32, 113–135.
- Sture Holm. 1979. [A simple sequentially rejective multiple test procedure](#). *Scand. Stat. Theory Appl.*, 6(2):65–70.
- Daniel Kahneman and Amos Tversky. 2013. Choices, values, and frames. In *Handbook of the fundamentals of financial decision making: Part I*, pages 269–278. World Scientific.
- Robin Lakoff. 1975. Linguistic theory and the real world 1. *Language Learning*, 25(2):309–338.
- Campbell Leaper and Melanie M Ayres. 2007. A meta-analytic review of gender variations in adults’ language use: Talkativeness, affiliative speech, and assertive speech. *Personality and Social Psychology Review*, 11(4):328–363.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- S Manikandan. 2020. A modern concept of blind hiring: Its importance and benefits. *Journal of Business and Management*, 22(5):60–62.
- Corinne A Moss-Racusin, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman. 2012. Science faculty’s subtle gender biases favor male students. *Proceedings of the national academy of sciences*, 109(41):16474–16479.
- D Neumark. 2010. Detecting discrimination in audit and correspondence studies (nber working paper series no. 16448).
- Matthew L Newman, Carla J Groom, Lori D Handelman, and James W Pennebaker. 2008. Gender differences in language use: An analysis of 14,000 text samples. *Discourse processes*, 45(3):211–236.
- Sheilla Njoto, Marc Cheong, Reeva Lederman, Aidan McLoughney, Leah Ruppner, and Anthony Wirth. 2022. Gender bias in AI recruitment systems: A sociological- and data science-based case study. To appear in: *Proceedings of the IEEE International Symposium on Technology and Society 2022*.
- Prasanna Parasurama and João Sedoc. 2022. [Gendered language in resumes and its implications for algorithmic bias in hiring](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 74–74, Seattle, Washington. Association for Computational Linguistics.
- Kim Parker. 2021. What’s behind the growing gap between men and women in college completion? Technical report, Pew Research Center, <https://www.pewresearch.org/fact-tank/2021/11/08/whats-behind-the-growing-gap-between-men-and-women-in-college-completion/>.
- Julie E. Phelan, Corinne A. Moss-Racusin, and Laurie A. Rudman. 2008. [Competent yet out in the cold: Shifting criteria for hiring reflect backlash toward agentic women](#). *Psychology of Women Quarterly*, 32(4):406–413.
- Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. [Connotation frames: A data-driven investigation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 311–321, Berlin, Germany. Association for Computational Linguistics.
- Peter A Riach and Judith Rich. 2006. An experimental investigation of sexual discrimination in hiring in the english labor market. *The BE Journal of Economic Analysis & Policy*, 6(2).
- Laurie A Rudman. 1998. Self-promotion as a risk factor for women: the costs and benefits of counterstereotypical impression management. *Journal of personality and social psychology*, 74(3):629.
- Laurie A Rudman and Peter Glick. 2001. [Prescriptive gender stereotypes and backlash toward agentic women](#). *J. Soc. Issues*, 57(4):743–762.
- Maarten Sap, Marcella Cindy Prasettio, Ari Holtzman, Hannah Rashkin, and Yejin Choi. 2017. [Connotation frames of power and agency in modern films](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2329–2334, Copenhagen, Denmark. Association for Computational Linguistics.
- Rhea E Steinpreis, Katie A Anders, and Dawn Ritzke. 1999. The impact of gender on the review of the curricula vitae of job applicants and tenure candidates: A national empirical study. *Sex roles*, 41(7):509–528.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. [Mitigating gender bias in natural language processing: Literature review](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1630–1640, Florence, Italy. Association for Computational Linguistics.
- Mary Talbot. 2019. *Language and gender*. John Wiley & Sons.

A Prolific consent form

Consent Form

I consent to participate in this project, the details of which have been explained to me.

I understand that the purpose of this research is to investigate how to produce a CV that increases success in recruitment.

I understand that my participation in this project is for research purposes only.

I acknowledge that the possible effects of participating in this research project have been explained to my satisfaction.

In this project, I will be required to draft a CV to apply for a promotion.

I understand that my participation is voluntary and that I am free to withdraw from this project anytime without explanation or prejudice and to withdraw any unprocessed data that I have provided.

I understand that the data from this research will be stored at the University of Melbourne and will be destroyed after 5 years.

I have been informed that the confidentiality of the information I provide will be safeguarded and subject to any legal requirements; my data will be password protected and accessible only by the university-approved researchers.

I understand that all data recorded is anonymized.

I understand that after I sign and return this consent form, it will be retained by the researcher.

By clicking the checkbox below, you are signing the Consent Form.

B Occupations

The occupational code was taken from the General Social Survey,⁹ as the Standard Occupational Classification System widely used in English-speaking official surveys. categorization Table 3 lists all occupations in our data set, together with the total number of CVs and gender ratio. "Other" occupations arise from free-text entries. Some examples include dog caretaker, musician, transport and logistics manager, pilot, automotive designer, among others. For reference, we also include gender ratios from the 2021 US Labor Statistics in the right-most column of Table 3. In general, the gender skew in our data set agrees with the US statistics, with some deviations expected given our limited sample.

⁹<https://gss.norc.org/>

C Gender-associated word choice

Table 4 shows the full results of gender-specific top 1% TFIDF terms per occupation (top), and overall across our whole CV data set (bottom).

D Classifier features

Table 5 lists the most predictive features for men (top) and women (bottom) as learnt by the linear SVC with L1 regularization when applied to the occupation-wise gender-balanced CV data set.

occupation	Man	Woman	% F	Total	% F USLS
Business and financial operations occupations	117	85	0.42	202	0.55
Computer and mathematical occupations	152	49	0.24	201	0.26
Educational instruction and library occupations	64	117	0.65	181	0.74
Healthcare support occupations	58	118	0.67	176	0.85
Management occupations	101	58	0.36	159	0.52
Sales and related occupations	65	70	0.52	135	0.62
Arts, design, entertainment, sports, and media occupations	58	64	0.52	122	0.50
Office and administrative support occupations	43	78	0.64	121	0.72
Other	61	53	0.46	114	–
Life, physical, and social science occupations	33	62	0.65	95	0.61
Architecture and engineering occupations	56	18	0.24	74	0.24
Community and social service occupations	15	44	0.75	59	0.72
Food preparation and serving related occupations	20	33	0.62	53	0.59
Legal occupations	28	22	0.44	50	0.42
Personal care and service occupations	8	21	0.72	29	0.72
Protective service occupations	5	5	0.5	10	0.50
Farming, fishing, and forestry occupations	2	6	0.75	8	0.75
Total	886	903	0.5	1789	0.53

Table 3: Occupations with number of CVs and proportion of female participants (%F) in our CV data set. The occupations included in our analyses in Appendix D are bold-faced. % F USLS are the official percentages of female employees per occupation taken from the US Labor Statistics (2021).

Gender balance	Occupation	Terms
Gender-balanced	Business and financial operations occupations	$M \setminus F$: insurance, sale, analysis, computer, word, lead, knowledge, master, data, product, time, operation, tax $F \setminus M$: student, august, accountant, art, high, information, public, use, training, software, development, research, program
	Sales and related occupations	$M \setminus F$: major, proficient, problem, engineering, june, computer, account, issue, lead, use, goal, responsibility, technical, develop, representative, operation, support, strategy $F \setminus M$: position, degree, industry, august, english, excel, excellent, leadership, strong, student, art, medium, create, study, maintain, good, employee, communication
Men-dominated	Computer and mathematical occupations	$M \setminus F$: gpa, application, database, office, window, microsoft, high, server, issue, engineer, web, security, include, developer, product, develop, network, college $F \setminus M$: student, art, analysis, lead, java, ms, python, course, sale, master, social, communication, spanish, css, time, html, graduate, ability, maintain, research
	Management occupations	$M \setminus F$: california, excel, leadership, engineering, ability, computer, company, information, planning, marketing, software, technology, product, time, administration, degree, support, design $F \setminus M$: staff, position, member, community, diploma, medium, account, research, health, hi, job, master, study, public, social, create, graduate, proficient
Women-dominated	Educational instruction and library occupations	$M \setminus F$: project, level, office, datum, <u>microsoft</u> , computer, <u>lead</u> , new, <u>software</u> , technology, lesson, business, develop, class, proficient $F \setminus M$: gpa, degree, library, member, <u>child</u> , community, award, create, elementary, honor, psychology, development, <u>social</u> , <u>help</u> , write
	Healthcare support occupations	$M \setminus F$: physician, various, equipment, project, member, volunteer, community, emergency, information, manage, manager, department, able, ensure, american $F \setminus M$: gpa, august, assistant, aid, problem, assist, june, art, microsoft, customer, paste, time, cpr, therapy, role
Overall: across all occupations		$M \setminus F$: web, server, improve, production, good, day, engineering, tool, build, policy, degree, check, increase, test, technology, meet, senior, master, material, security, engineer, control, procedure, sql, solution, point, administration, performance, quality, database, network, equipment, data, strategy, user, testing $F \setminus M$: strong, psychology, able, conduct, word, care, position, organize, prepare, art, document, intern, national, coordinate, resource, online, record, schedule, teacher, space, course, gpa, teach, english, child, current, event, class, store, volunteer, september, meeting, honor, individual, content, study, phone

Table 4: Qualitative analyses of mutually-exclusive terms within the top 1% of unigrams, by tf-idf ranking: (top) stratified across occupations; and (bottom) across all occupations. $M \setminus F$ denotes the set of terms in the top 1% of male CV unigrams, which are not in the corresponding top 1% of female CV unigrams. Conversely, $F \setminus M$ denotes the set of terms found in the top 1% of female CV unigrams, which are not in the top 1% of male CV unigrams. The underlined examples are discussed in Section 3.

male: level, indianapolis, instruct, reduce, th, clinical, troubleshoot, large, part, culinary, engineer, supervise, repair, improvement, shipping, technology, multiple, tool, business, regulatory

female: answer, attitude, reference, file, document, create, assist, role, colorado, children, child, coordinator, know, gain, interview, receivable, honors, woman, media

male -NE: machine, instruct, profit, clinical, also, supervise, observe, part, technology, leader, reduce, reduction, basic, instructor, opportunity, hold, review, people, regard, electrical
female -NE: check, attitude, document, core, medium, media, woman, answer, content, present, child, create, speaking, claim, resource, file, assist, resume, plan

Table 5: Most predictive features as learned by the binary SVC for gender prediction when trained on the Balanced data set full (top) or with NEs redacted (bottom).