

Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval

Kyra Wilson, Aylin Caliskan

University of Washington Information School
 Seattle, Washington, USA
 {kywi, aylin}@uw.edu

Abstract

Artificial intelligence (AI) hiring tools have revolutionized resume screening, and large language models (LLMs) have the potential to do the same. However, given the biases which are embedded within LLMs, it is unclear whether they can be used in this scenario without disadvantaging groups based on their protected attributes. In this work, we investigate the possibilities of using LLMs in a resume screening setting via a document retrieval framework that simulates job candidate selection. Using that framework, we then perform a resume audit study to determine whether a selection of Massive Text Embedding (MTE) models are biased in resume screening scenarios. We simulate this for nine occupations, using a collection of over 500 publicly available resumes and 500 job descriptions. We find that the MTEs are biased, significantly favoring White-associated names in 85.1% of cases and female-associated names in only 11.1% of cases, with a minority of cases showing no statistically significant differences. Further analyses show that Black males are disadvantaged in up to 100% of cases, replicating real-world patterns of bias in employment settings, and validate three hypotheses of intersectionality. We also find an impact of document length as well as the corpus frequency of names in the selection of resumes. These findings have implications for widely used AI tools that are automating employment, fairness, and tech policy.

Introduction

One of the widespread practical applications of artificial intelligence (AI) tools has been their use in hiring processes. It is estimated that 99% of Fortune 500 companies are already using some sort of AI assistance when making hiring decisions (Schellmann 2024), due to their potential to increase recruitment quality and efficiency (Chen 2023). This could potentially alleviate discrimination based on people’s unconscious biases or stereotypes, such as associating Black male job seekers with criminals (Pager 2003) or female job seekers with lower productivity due to motherhood (González, Cortina, and Rodríguez 2019). However, many AI hiring tools do still exhibit biased outcomes, such as a resume screening tool developed at Amazon which had to be scrapped when it was revealed that it unfairly discriminated against women (Dastin 2018).

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Large language models (LLMs), which are trained on a general corpus of language data rather than data specific to hiring tasks, also have the potential to be used in these scenarios. In fact, resume screening tasks have already been observed in interactions between users and ChatGPT (Ouyang et al. 2023). Additionally, the accessibility of these models (both in terms of cost and user interfaces) lend them well to adoption by companies which have either not yet incorporated AI assistance into their hiring pipeline due to cost or technological complexity or to replace models that are currently in use and whose biases may be better understood.

With the potential for increased use of LLMs in hiring, it is essential to document the extent to which they exhibit biases against particular social groups. In the US, it is illegal to make hiring decisions on the basis of race, color, religion, sex (including gender identity, sexual orientation, and pregnancy), national origin, age (40 or older), disability or genetic information. A number of these biases have been documented in LLMs already, including gender, race, religion, and disability biases, as well as their intersections (Kotek, Dockum, and Sun 2023; Narayanan Venkit 2023; Kirk et al. 2021; Guo and Caliskan 2021). Therefore, it is crucial to investigate whether these models exhibit discriminatory biases related to protected attributes such as race and gender or their intersections in order to evaluate how they can be used for resume screening tasks.

Additionally, it is also essential to investigate low-level textual features such as term frequencies and document lengths which, although not direct signals of social groups, play a significant role in the performance and outputs of language models (Jones and Steinhardt 2022; Anil et al. 2022; Wolfe and Caliskan 2021). These features can vary widely across documents such as resumes, and evaluating their relationship to social group outcomes in automated resume screening is essential in order to accurately represent real-world usage.

In this study, we formulate resume screening as a practical zero-shot document retrieval task. Using this approach, we seek to address the following questions:

- **RQ1:** Are identical resumes with different race (Black vs. White) or gender (male vs. female) signals selected at equal rates when using three state-of-the-art LLMs for resume screening via a practical retrieval task?
- **RQ2:** Are identical resumes with different intersectional

group signals selected at equal rates when using three language models for the same resume screening task?

- **RQ3:** How do the features of race and gender signals such as name frequency and resume length impact screening outcomes?

We introduce a simulation for LLMs' usage as resume screening tools and analyze outcomes with respect to race and gender. In a hiring decision pipeline, resume screening is generally considered to be the second stage, after sourcing potential candidates, and prior to interviewing and selecting a candidate (Bogen and Rieke 2018). LLM assistance is particularly useful for this stage because it involves the analysis of many text documents (resumes) to identify those that are most relevant relative to a particular job description.

We investigate LLM-mediated resume screening using publicly-available English datasets of over 500 resumes and 500 job descriptions that are collected from real-world examples of these documents. They span nine occupations (Chief Executive, Marketing and Sales Manager, Miscellaneous Manager, Human Resources Worker, Accountant and Auditor, Miscellaneous Engineer, Secondary School Teacher, Designer, and Miscellaneous Sales and Related Worker). Resume screening is analyzed within three different Massive Text Embedding (MTE) models, a special class of LLMs which are trained for representational tasks such as document retrieval, classification, and clustering after pre-training on general language corpora. While many studies have characterized the biases of foundation or instruction-tuned LLMs, very few have investigated the biases of MTEs or their use in resume screening, adding further novelty and importance to this study.

To audit the MTE models for biases in resume screening scenarios, we augment resumes with 120 frequency-controlled names that are associated with White, Black, male, and female identities. Using document retrieval and selecting subsets of resumes which are most similar to job descriptions, we are able to analyze whether outcomes differ across identity groups. The code and data produced by this research are available at <https://github.com/kyrawilson/Resume-Screening-Bias>. We make the following knowledge contributions:

1. In all of the models, we find that resumes that belong to the same occupation category as a given job description have significantly higher cosine similarities than resumes that belong to a different occupation category (0.0437 higher on average), justifying the use of cosine similarity to determine retrieval in order to analyze an initial stage of resume screening where a pool of the most relevant candidates are identified for further evaluation by a hiring professional.
2. Using more than three million comparisons between resumes and job descriptions, we find that resumes with White or male names are preferred to those with Black or female names in up to 85.1% of cases.
3. Intersectional comparisons reveal resumes that contain Black male names are highly disfavored in resume screening, with other groups being preferred in up to 100% of cases. Gender differences are driven largely by

disparate preferences of Black females over Black males, as White males and White females have much smaller selection rate differences. These findings validate three hypotheses of intersectionality.

4. Features such as resume length and name frequency significantly impact bias measurements in LLM resume audits, such that increasing the ratio of signals that are proxies to race or gender information in a document by decreasing its length can increase the number of biased outcomes by 22.2%, and changing frequency matching strategies can alter whether Black names or White names are favored in a majority of cases.

Finally, we discuss how the resume screening patterns found in LLMs echo societal patterns of hiring discrimination and how the particular features of resumes such as names should be considered when using LLMs for resume screening.

Related Work

There has been limited work addressing and documenting the potential risks of using LLMs for hiring decisions, despite the documentation of biases in real-world resume screening scenarios and in AI models specially trained for hiring tasks. In economic labor market studies, bias in resume screening is usually investigated through resume audit or correspondence studies (Baert 2018). In these experiments, artificial resumes which differ only on some protected attribute are created and then sent in response to real job postings. Hiring procedures are determined to be discriminatory if the response rates differ significantly between groups that vary on the key dimension. This paradigm has been used to identify bias related to a number of protected attributes such as race (Bertrand and Mullainathan 2004), queerness (Mishel 2016), religion (Wright et al. 2013), disability (Hipes et al. 2016), and age (Neumark, Burn, and Button 2016, 2019; Lahey 2008).

To date, there are no external gender or racial bias audits of AI-mediated resume screening tools, the majority of which are typically closed-source, proprietary, and not accessible for external review (Li and Goel 2024). Limited work has addressed this issue by reviewing publicly available statements and model descriptions (Raghavan et al. 2020; Sánchez-Monedero, Dencik, and Edwards 2020), finding that the majority of vendors do not make explicit statements regarding their models' compliance with anti-discrimination law, and those that do are typically only within a US context. Wilson et al. (2021), using a cooperative audit, found that the system of interest did not exhibit adverse impact, but key assumptions make this result difficult to generalize without additional testing. A final external audit found that closed-source models are often unstable, but they did not investigate any protected characteristics (Rhea et al. 2022).

Only very preliminary work has been done to investigate LLMs used for resume screening. A team of Bloomberg reporters investigated OpenAI's GPT-3.5 and GPT-4 and found that Black women were only ranked as top candidates for software engineering roles in 11% of tests, and Hispanic women were twice as likely as men to be ranked as top

candidates for human resources workers. The only career in which no group was disadvantaged was retail workers (Yin, Alba, and Nicoletti 2024). Another study investigating OpenAI’s ChatGPT found that resumes which mentioned disability in the context of an award were only ranked highest in 25% of cases (Glazko et al. 2024). While both studies demonstrate biased outcomes when using LLMs or chatbots as resume screeners, the models investigated were all “black boxes,” meaning the analysis was limited to model outputs only and could not investigate internal representations. Additionally, researchers did not rigorously investigate low-level document features such as term frequency or length in their studies, which are related to model biases (Esiohu et al. 2023).

Another limitation of studies investigating chatbots’ use as resume screeners is that they are less transparent and rely on the model’s generative capabilities for resume screening. Language generation is relatively computationally expensive as new tokens are produced iteratively, and outputs are also highly sensitive to features of the prompt which are unrelated to the task itself (Sclar et al. 2023). An alternative is to use embeddings of documents directly, which are less computationally intensive to process although still potentially sensitive to spurious prompt features.

Data

The present study examines race and gender bias in resume screening using a corpus of resumes and job descriptions as well as three MTEs (illustrated in Figure 1) with a Mistral-7B-like architecture (Jiang et al. 2023). Mistral-7B is based on the transformer architecture (Vaswani et al. 2017), but it innovates using grouped-query attention and sliding-window attention to decrease computational costs while also being able to process long sequences more effectively, a key innovation for the resume screening setting where long documents are common.

Massive Text Embedding Models

We select three high-performing MTEs to illustrate the potential range in outcomes from a set of architecturally similar models. The chosen models were fine-tuned for representational tasks (including document retrieval, classification, and clustering) on either Mistral-7B-v0.1, a foundation model whose learning objective is solely next-token prediction, or on models which were themselves fine-tuned on Mistral-7B-v0.1. In order to achieve high performance on these tasks, LLMs must undergo an additional round of training (*fine-tuning*) in which a specialized loss function is used to update model parameters. Typically, the loss function used for representational tasks is a contrastive loss (CL), in which embeddings of examples that are positive examples of a query or label are brought closer together spatially, while those that are negative examples are pushed further apart (Jones and Steinhardt 2022).

E5-mistral-7b-instruct (e5) (Wang et al. 2023) is trained on Mistral-7B-v0.1 and fine-tuned using using a CL objective with natural and synthetically generated training data for retrieval tasks. GritLM-7B (*GritLM*) (Muennighoff et al.

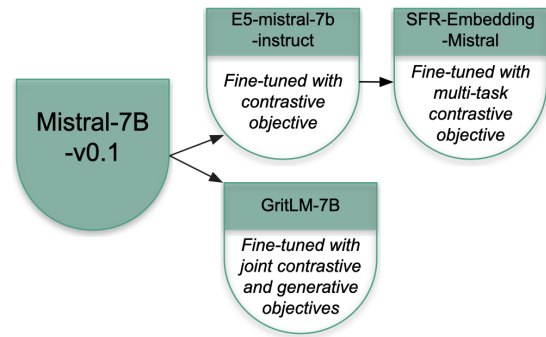


Figure 1: Relation of the three models investigated to a pre-trained LLM. Arrows between models indicate additional fine-tuning steps.

2024) is also fine-tuned on top of Mistral-7B-v0.1; however, it has no synthetic data in its training set and uses a joint objective with both CL and next-token prediction in order to train for both representational and generative functions. Finally, SFR-Embedding-Mistral (*SFR*) (Meng et al. 2024) is fine-tuned with CL on top of e5, which itself is already a fine-tuned MTE. Its training data includes both retrieval tasks as well as classification and clustering tasks.

All three MTEs achieve state-of-the-art performance according to the Massive Text Embedding Benchmark (MTEB), which aims to quantify the performance of MTEs on a large set of representational tasks and datasets (Muennighoff et al. 2022). On document retrieval tasks, they are the highest performing open-source MTEs with at least one billion parameters¹.

Job Description and Resume Stimuli

In order to simulate resume screening, a selection of job descriptions as well as candidate resumes are necessary. While examples of job descriptions are widely available online, they are frequently removed or altered as positions fill or requirements change. This makes reproducing research using these descriptions difficult. Additionally, resumes contain a wealth of sensitive and private information, so can be difficult to both collect and disseminate. To ensure reproducibility, we select two job description and resumes datasets which are publicly available online². Job descriptions are a free-text solicitation for an open position; resumes have an occupation title followed by free-text description of qualifications.

Because these datasets were not annotated with the same occupation classification schemes, it was necessary to standardize the job titles associated with each document in the dataset. The US government developed the Standard Occupational Classification (SOC) system in order to classify and collect data about particular occupations. Using this scheme,

¹As of April 2024.

²Resumes were collected from <https://www.kaggle.com/datasets/snehaanbhawal/resume-dataset>. Job descriptions were collected from <https://www.kaggle.com/datasets/marcocavaco/scraped-job-descriptions>.

Broad Occupation	% Women	% White	% Black	Total Workers (Thousands)	# of Resumes	# of Descriptions
Chief Executives	30.6	85.8	5.2	1,780	49	33
Marketing and Sales Managers	47.7	86.9	5.0	1,136	51	81
Miscellaneous Managers	37.5	80.7	9.2	5,666	112	92
Human Resources Workers	76.5	74.9	14.7	980	22	27
Accountants and Auditors	57.0	73.4	11.9	1,624	133	99
Miscellaneous Engineers	15.4	72.4	5.9	669	28	112
Secondary School Teachers	56.9	87.8	6.1	944	20	52
Designers	52.0	80.5	6.6	862	87	55
Miscellaneous Sales and Related Workers	56.1	81.0	12.0	339	52	20

Table 1: Resume screening was investigated for nine occupations. These are presented with corresponding US population statistics for percentage of women workers, White workers, Black workers, and total number of workers as well as the number of documents corresponding to each occupation category from the resume and job description datasets after filtering.

six-digit codes were assigned for each job description and resume using the NIOSH Industry and Occupation Computerized Coding System (NIOCCS)³ from the US government. This tool was developed to programatically assign SOC codes to free text job titles and descriptions⁴. Code assignments with less than 60% confidence were eliminated from that dataset. This threshold was chosen to maximize code confidence while also maintaining a large dataset with a range of occupations represented.

Using the broad occupation (first five digits of the SOC code) identified for each resume and job description, we filtered the dataset to remove duplicates and keep only documents belonging to occupation categories which had at least 20 resumes and 20 job descriptions, ensuring a large pool of documents for accurately simulating resume screening and achieving necessary statistical power. After this, the dataset contained 554 resumes and 571 job descriptions across nine occupations (Chief Executive, Marketing and Sales Manager, Miscellaneous Manager, Human Resources Worker, Accountant and Auditor, Miscellaneous Engineer, Secondary School Teacher, Designer, and Miscellaneous Sales and Related Worker). A complete listing of the categories and associated employment statistics from the Bureau of Labor Statistics (BLS) can be seen in Table 1.

Approach

The present study describes a generalizable and scalable approach to resume screening based on document retrieval.

³<https://csams.cdc.gov/nioccs/>

⁴Schmitz, Forst et al. (2016) investigated the NIOCCS with respect to manual coding and found that the tool had “fair” to “good” accuracy when assigning codes with up to four digits. Although this is lower than what is reported by NIOSH, we find that this tool still has superior performance on our sample of resumes and job descriptions in comparison to other SOC coding tools.

This involves using MTEs to create embeddings for job descriptions and resumes, and then using a simple cosine similarity comparison to capture which resumes are most similar to a given job description. A summary of this approach is given in Figure 2. A chi-square test is then used to determine whether the most similar resumes are distributed equally amongst relevant groups, or whether certain groups are favored over others, indicating bias. Chi-square tests require a minimum of five observed values for valid population estimates (Franke, Ho, and Christie 2012); our dataset far exceeds that, with at least 160 resume documents for every bias test, demonstrating the legitimacy of the results at scale.

Task-Augmented Job Descriptions

In document retrieval with MTEs, query texts are encoded with an additional instruction describing the particular setting.⁵ We created a set of 10 instructions, shown in Table ??, to be encoded with job descriptions according to templates specific to each model. We used ChatGPT in this procedure to develop alternatives for the phrases “job description” and “resume.”

Name-Augmented Resumes

To measure bias in resume screening, resumes were augmented with a name, comprised of a variable first name and constant last name, by prepending the complete name to the beginning of the document. *Williams* was selected as a last name because it is both frequent (third most common name in the US) and approximately equally likely to be used either by a Black or White person (47.68% vs. 45.75%) (Census 2024). The last name was kept constant across all resumes

⁵The text formatting used to encode instructions and query texts varies between MTEs. For each model, we followed the recommended structure as described in that model’s documentation for these experiments.

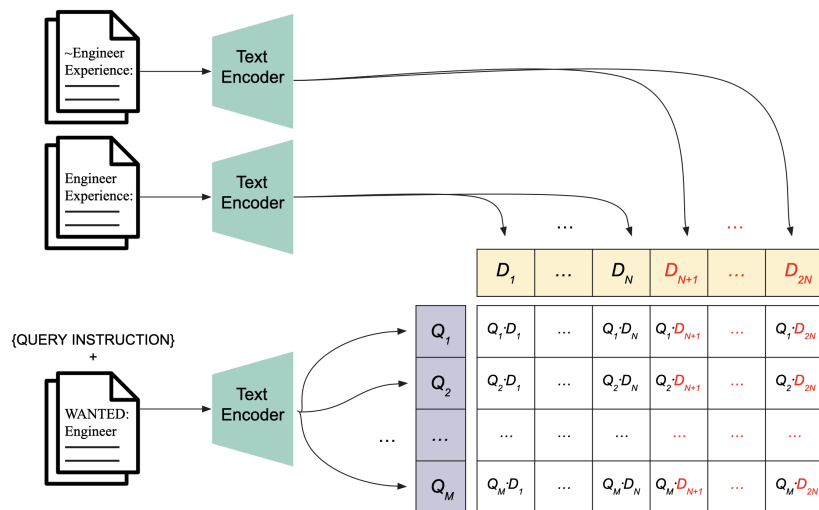


Figure 2: Illustration of the resume screening as document retrieval framework. Task instructions are appended to job descriptions and treated as queries, while resumes are treated as documents. The cosine similarity between queries and documents estimates the relevance of a resume to a particular job description.

in order to maximize experimental control and document realism while also minimizing required computation.

We use the name database introduced in Elder and Hayes (2023) to select names associated with one of four groups: Black males, Black females, White males, or White females. Of these, the Black male group had the fewest potential names, and the top 20 most distinctive⁶ names (33% of all Black male names in the database) were chosen for use in resume augmentation.

An equal number of names corresponding to other groups were then selected in order to closely match or be proportional to the corpus frequencies of the Black male names. Corpus frequencies were determined using infini-gram (Liu et al. 2024), a tool that facilitates n-gram searches for arbitrarily large corpora, and the DOLMA corpus (Soldaini et al. 2024)⁷.

The first set of names was selected in order to be proportional to the relative population differences between Black and White people in the US, replicating the distribution of names that would likely be seen in real-world resume screening. According to 2023 US Census estimates, those who identify as White alone comprise 75.5% of the US population, while those who identify as Black alone comprise 13.6%. Accordingly, we selected White male and female names which were approximately 5.5 times more frequent in the corpus than corresponding Black male names, and Black

⁶Distinctiveness was measured via the difference in ratings from one to five between the most likely racial group versus the second most likely racial group, with all names having a score of at least 0.66.

⁷Although Mistral model weights are available publicly, the training dataset is proprietary. The DOLMA frequencies are meant only to approximate the frequencies of names in the corpus used to train Mistral. DOLMA contains 3.1 trillion tokens and is currently the largest available to search using infini-gram.

female names which were approximately equally frequent to Black male names. We created an additional set of 20 White male and 20 White female names which were also approximately equally frequent to Black male names for supplemental investigation of frequency effects. A sample of first names from both the proportional frequency-matched and exact frequency-matched sets can be seen in Table 2 with a full list available in the Appendix.⁸

Resume Screening

Zero-shot dense retrieval, which uses contextualized embeddings to compare documents rather than exact term matches, provides a natural analog for resume screening. In the initial stages of retrieval, relevance scores computed from text embeddings are used to select a set of documents from a large corpus that best match a user query. Cosine similarity is commonly used as a relevance metric (Zhao et al. 2024). Similarly for resume screening, resumes r which are similar to a job description d can be identified via their respective embeddings v_r and v_d using the equation in (1). Furthermore, using a retrieval approach for resume screening allows for the direct analysis of textual embeddings to determine whether the representations are potentially biased in a way that could influence model outputs. If the resumes which are most similar to a particular job description consistently belong to a certain group, this is evidence that the representations are biased in favor of that group.

$$\text{sim}(r, d) = \langle v_r, v_d \rangle \quad (1)$$

Embeddings for document retrieval were generated using the MTEs in Figure 1. Texts were truncated due to computational limitations, and 1,300 tokens was chosen as a

⁸A long-form version of this paper with Appendix is available at <https://arxiv.org/abs/2407.20371>.

Name	Lg. Freq.	Distinct	Group
Kenya	16.87	1.80	BF
Ebony	15.18	1.29	BF
Latrice	10.86	0.71	BF
Latisha	10.77	1.57	BF
Jackson	17.71	0.67	BM
Abdul	16.18	0.82	BM
Demetrius	13.14	0.89	BM
Dewayne	12.01	0.97	BM
May	19.45	0.66	WF
Hope	17.86	1.07	WF
Stacy	14.85	1.37	WF
Kristine	13.73	1.30	WF
John	19.22	0.89	WM
Joe	17.90	1.14	WM
Stevie	14.88	0.66	WM
Huey	13.73	0.69	WM
Virginia	17.84	1.05	=WF
Katie	16.18	1.66	=WF
Aileen	13.12	0.87	=WF
Rebeca	11.97	1.30	=WF
Daniel	17.71	1.06	=WM
Spencer	16.18	0.91	=WM
Bennie	12.95	0.70	=WM
Wilbert	12.317	0.87	=WM

Table 2: The two most and least frequent first names used for each intersectional group along with their corpus frequencies and racial distinctiveness scores. = indicates White names which are matched exactly to the frequencies of the Black names rather than using names whose frequencies are proportional to US population differences.

length which captured the majority of resume content while still being computationally feasible. A summary of unaltered document lengths is available in the appendix. Document embeddings were extracted from the last hidden state of a model and normalized before computing their cosine similarity. Cosine similarity scores were averaged over task instructions, so the updated similarity computation is as in Equation (2), where t corresponds to the index of the task instructions in Table ?? used to form job description embeddings. For completeness, results for individual task instructions are also provided in the Appendix.

$$sim(r, d) = \frac{1}{10} \sum_{t=1}^{10} \langle v_r, v_{d_t} \rangle \quad (2)$$

To simulate candidate selection, we select a percentage of the most similar of resumes for each job description for further analysis. A chi-square test is used to determine whether the selected resumes are distributed uniformly amongst relevant groups or whether particular groups are represented at significantly higher rates than others, indicating bias in resume screening outcomes. Results for resume screening outcomes are presented primarily in terms of difference in selection rates; intermediate cosine similarity results are provided in the Appendix.

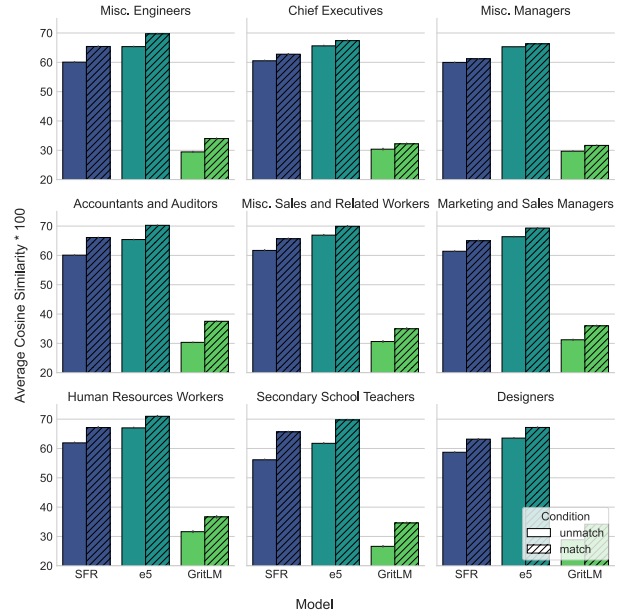


Figure 3: For each occupation and model, cosine similarities are significantly higher ($p < 0.001$) for resumes which belong to the same occupation as the job description (*match*) than those that belong to different occupations (*unmatch*), indicating the success of the document retrieval for resume screening framework.

Experiments

Detailed information is provided for four experiments. Experiment 1 evaluates the document retrieval for resume screening framework by comparing the similarity of resumes which belong to the same occupation category as a given job description to those that belong to different categories. Experiment 2 and 3 investigate bias in resume screening, first using gender and race categories separately, then investigating intersectional identities. Finally, Experiment 4 examines the effect of low-level features such as name frequency and resume length on bias measurements. In all experiments, the initial pools of resumes are balanced with respect to identity groups, and the expected outcome is that all groups should be represented uniformly if the MTEs are unbiased. Any significant deviations from this represent biased outcomes against particular identity groups in resume screening.

Evaluating Retrieval for Resume Screening

Cosine similarity of embeddings for resumes without names and job description embeddings was calculated for both resumes whose occupation category corresponded to the job description (*matched*) and those which did not (*unmatched*), simulating the initial stage of screening resumes for relevance. The cosine similarity scores of these two groups were compared to verify that the document retrieval approach is suitable for resume screening with LLMs.

Evaluating Race and Gender Bias

Gender and race groups were formed by combining names with population proportional frequencies from the four intersectional groups into four groups corresponding to only one race or gender identity (Black, White, male, or female). Each group was comprised of 40 names. Embeddings for job descriptions and name-augmented resumes were created using the three MTE models and cosine similarities were computed.

For each model and occupation, we performed a bias test by selecting the top 10% of most similar resumes for every job description and determining whether race or gender groups were represented at significantly higher rates. At this threshold, a minimum of 160 resumes were selected for each job description, and a total of 27 bias tests were conducted for both gender and race.

Evaluating Intersectional Bias

Using the 20 names with population proportional frequencies from each intersectional group (Black female, Black male, White female, White male), we repeated the embedding procedures, selection of top 10% of resumes, and 27 chi-square bias tests from Experiment 2 for each pair of intersectional identities, excluding those in which no race or gender dimension was shared.

Evaluating Effects of Length and Frequency

An additional set of name-augmented resumes was created in which the document contained a name and occupation title only (*title-only*) in contrast to the original name-augmented resumes which contained a name, occupation title, and additional content (*full-length*). Cosine similarities and bias tests were repeated as described above, and the results of using *title-only* versus *full-length* resumes were compared for non-intersectional race and gender categories.

Finally, the effect of using raw frequency matched versus population proportional frequency matched names was investigated by repeating the cosine similarity and bias tests as described above for resumes augmented with both sets of names, and comparing the results for non-intersectional race and gender categories.

Results

First, evidence indicates successful resume screening via retrieval. Additional evidence indicates underlying biases favoring resumes with White or male names when race and gender are analyzed independently. Intersectional analyses indicate this bias is strongest against resumes with Black female or male names. Finally, attributes such as name frequency and resume length also effect the similarity of resumes to job descriptions.

Verification of Retrieval

Resumes with no names belonging to the same occupation category as the job description have significantly higher cosine similarities ($p < 0.001$) than resumes with no names belonging to different occupation categories for every occu-

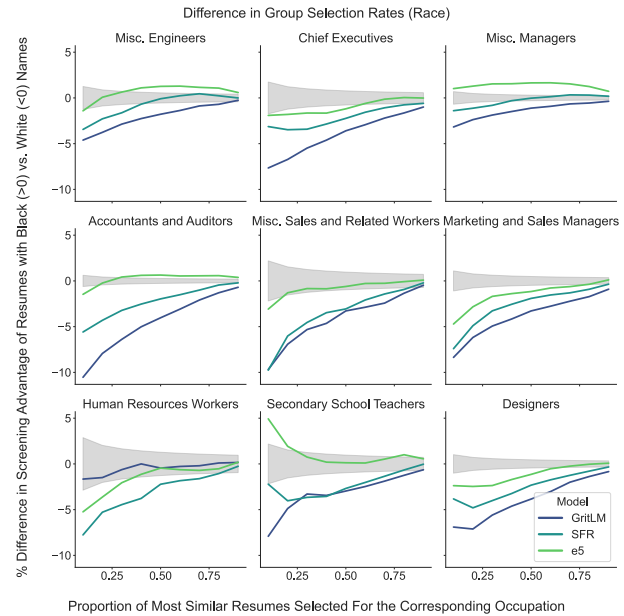


Figure 4: Resumes with White names are significantly preferred ($p < 0.05$) in 85.1% of tests; those with Black names are preferred in 8.6% of tests. Gray regions indicate disparities which are not significantly different from zero (6.3% of tests).

pation, as seen in Figure 3, verifying the use of document retrieval for resume screening.

White and Male Names are Preferred

The majority of experiments reveal a preference for White names over Black names. In 85.1% of 27 tests for racial bias in resume screening, White names were preferred, and only in 8.6% of tests were Black names preferred, as seen in Figure 4. Notably, the only cases in which Black-associated names were preferred at any threshold were results from the e5 model, suggesting that this MTE may be less racially biased than others for particular occupations.

While male names were also favored compared to female names in the majority of experiments, the disparities were less than those demonstrated using Black versus White names. In 51.9% of 27 tests for gender bias in resume screening, male names were preferred to female names, and only in 11.1% of tests were female names preferred, as seen in Figure 5. Again, all cases in which female names were preferred came from a single model, GritLM.

Nuances of Intersectional Identities

Intersectional comparisons reveal that the smallest disparities exist between White names of different genders. Comparisons between Black names of different genders or Black and White names of the same gender exhibit larger disparities. Comparisons between resumes with White male and White female names reveal significant differences ($p < 0.05$) between the two groups in only 44.4% of 27 tests, as shown

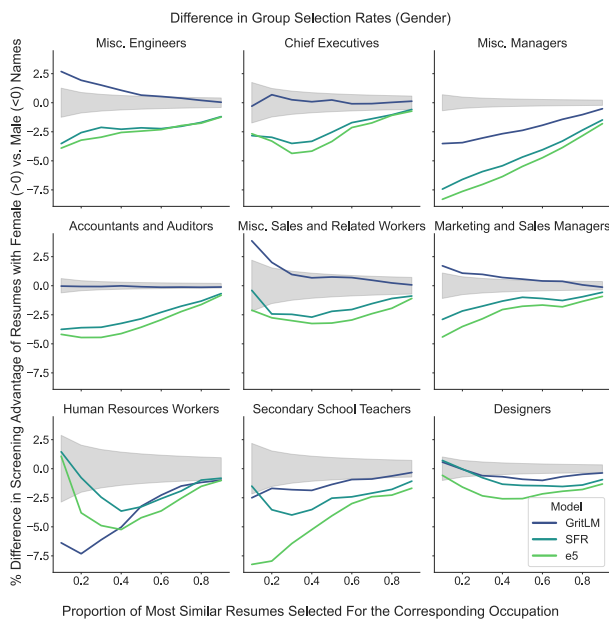


Figure 5: Resumes with male names are significantly preferred ($p < 0.05$) in 51.9% of tests; those with female names are preferred in 11.1% of tests. Gray regions indicate disparities which are not significantly different from zero (37% of tests).

in Figure 6. White males are only preferred in 18.5% of tests and White females are preferred in 25.9%. No model exhibits consistent behavior in preferring one group where other models do not.

Larger differences appear in comparisons including Black names. Tests of resumes with White female names versus Black female names show a statistically significant preference for the former in 48.1% of cases and the latter in only 25.9% of cases ($p < 0.05$), as shown in Figure 7. Conversely, tests of resumes with Black female names versus Black male names reveal a significant preference for the former in 66.7% of tests and the latter in only 14.8% of tests ($p < 0.05$), as shown in Figure 8. GritLM seems to exhibit the most consistent biases against resumes with Black female names, as this is the only model which always prefers resumes with White female names and has the highest rate of preference for Black male names as well. Finally, comparisons of resumes with Black male names to those with White male names reveal the largest biases, as in Figure 9. White male names are significantly preferred in all 27 tests ($p < 0.05$).

Shorter Resumes Result in More Bias

Overall, we find that resumes with names and title only (*title-only*) lead to more biased outcomes than those with names, titles, and content (*full-length*). For title-only race tests, significant differences are observed in 96.2% of 27 bias tests, compared to 93.7% of bias tests of full-length resumes ($p < 0.05$). Of these, resumes with White names were significantly preferred in 62.9% of tests; Black names were

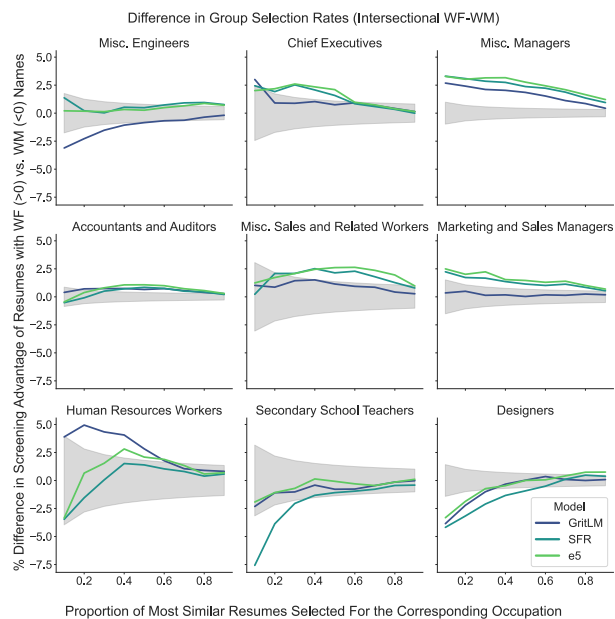


Figure 6: Resumes with White male names are preferred in 18.5% of tests; those with White female names are preferred in 25.9%. Gray regions indicate disparities which are not significantly different from zero (55.6% of tests).

preferred in 33.3%.

For title-only gender tests, significant differences are identified in 85.2% of cases, compared to 63% for full-length cases; this increase is entirely attributable to an increase of significant preferences for resumes with female names. Additionally, for both gender and race tests, the difference between group selection rates increases for title-only resumes. Detailed results for title-only resumes can be seen in the Appendix.

Frequency Effects Bias Measurements

Finally, the frequency of names also had a significant impact on outcomes. When names which had approximately equivalent frequencies in the DOLMA corpus were used, we find that resumes with Black names are preferred to those with White names in 51.9% of 27 tests, while those with White names are preferred to those with Black names in only 22.2% of tests. This is the reverse pattern than what is observed when names are selected based on population proportional frequencies, exemplifying model sensitivity to low level features such as name frequency. Results figures for these alternate resumes can be seen in the Appendix.

Discussion

The results of these experiments illuminate the potential for biased outcomes when using LLMs as resume screeners, as each of the three MTEs had outcomes which favor certain social groups over others. When analyzing race and gender independently, we find that the MTEs show an overall preference for resumes with White and male names, rather than

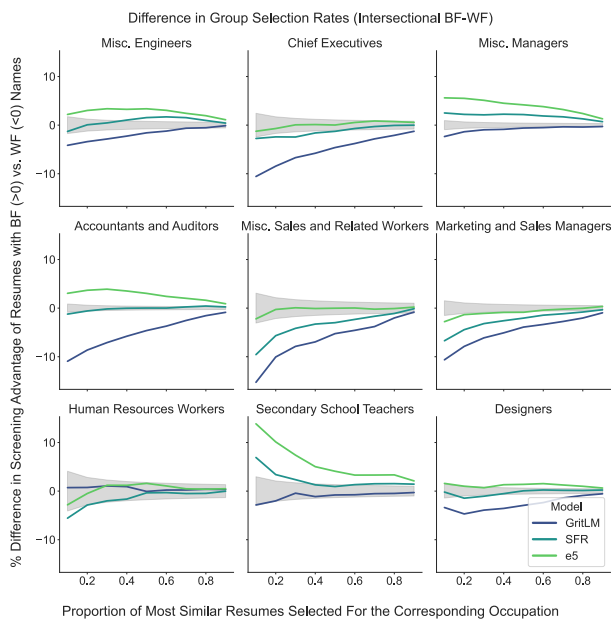


Figure 7: Resumes with White female names are preferred in 48.1% of tests; those with Black female names are preferred in 25.9%. Gray regions indicate disparities which are not significantly different from zero (26% of tests).

preferences that align with societal patterns. For example, resumes with Black names were preferred only for the occupations Miscellaneous Managers and Secondary School Teachers; resumes with female names were preferred for Miscellaneous Engineers, Miscellaneous Sales and Related Workers, and Marketing and Sales Managers. In neither of these cases are the occupations those that have the largest proportions of workers belonging to the preferred group according to the BLS, and further investigation is needed to explain these preferences.

The lack of correlation between population statistics and model preferences suggests that group disparities in resume screening are a consequence of default model preferences rather than occupational patterns learned during training. This contrasts to work which finds evidence of occupational biases in LLMs (Kotek, Dockum, and Sun 2023), and instead aligns with research suggesting that LLMs have an overall preference for certain social groups (Caliskan et al. 2022; Wolfe and Caliskan 2022; Ghosh and Caliskan 2023; Cheng, Durmus, and Jurafsky 2023). Specifically, masculine and White concepts seem to be treated as the “default” value by models with other identities diverging from this, rather than a set of equally distinct alternatives. In our experiments, this may be due to the fact that White and male defaults exist across many contexts, so occupation-based nuances may be weaker or lost all together.

Intersectional results, on the other hand, do correspond more strongly to real-world discrimination in resume screening. In this setting, Black males are consistently observed to be disadvantaged relative to all other Black and White workers (Pager 2003). Our results support this; resumes with

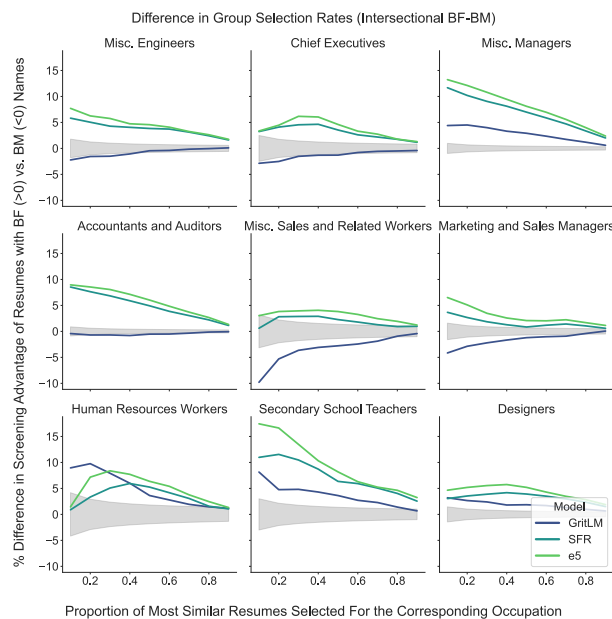


Figure 8: Resumes with Black female names are preferred in 66.7% of tests; those with Black male names are preferred in 14.8%. Gray regions indicate disparities which are not significantly different from zero (18.5% of tests).

Black male names are only preferred to Black female names and White male names in 14.8% and 0% of bias tests, respectively.

Furthermore, the results correspond to three additional hypotheses from Ghavami and Peplau (2013) regarding intersectionality. First, we observe that intersectional biases are not explained by the sum of race and gender biases alone. Second, the independent race bias tests align more closely with intersectional results between groups of males than groups of females. Finally, gender biases are more similar to biases between White males and females than Black males and females.

In addition to names affecting the resume screening outcomes, low-level features such as document length and name frequency also significantly altered outcomes. These manipulations both increased the amount of biased outcomes as well as which social groups were preferred, respectively. In a real-world resume screening scenario, it would be difficult to control these naturally varying features, which have an impact on LLM outcomes and may further disadvantage certain groups (Jones and Steinhardt 2022; Anil et al. 2022; Wolfe and Caliskan 2021).

While there are a number of factors contributing to biased outcomes in resume screening via LLMs, one naive approach to mitigation might be removing names from resumes altogether. However, resumes from real-world job seekers differ on many additional dimensions which can signal social group membership, including educational institutions, locations, and even lexical content choices. For example, Parasurama and Sedoc (2022) find that resumes written by women were more likely to use words like *cared* or *vol-*

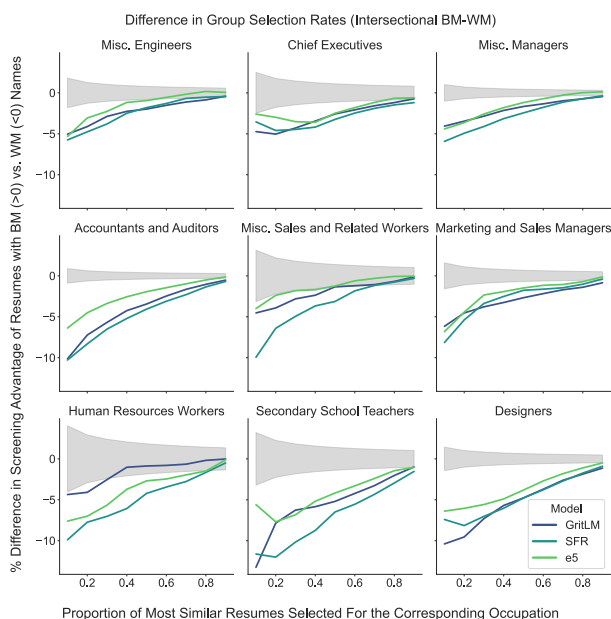


Figure 9: Resumes with White male names are preferred in 100% of tests; those with Black male names are preferred in 0%. Gray regions indicate disparities which are not significantly different from zero (0% of tests).

unteered, while men used words like *repaired* or *competed*, and these differences correlated with differences in hiring outcomes. While this study manipulated names alone due to their strong associations with certain protected characteristics, LLMs used for resume screening are likely sensitive to such signals stemming from additional sources as well.

Other approaches to bias mitigation have identified methods which can minimize nuanced social group signals from words other than names (Deshpande, Pan, and Foulds 2020; Parasurama and Sedoc 2022). Additionally, other methods such as debiasing embeddings or reranking documents have also been proposed to counteract biases in hiring and retrieval settings (Gerritse and de Vries 2020; Sundararaman and Subramanian 2022; Parasurama and Sedoc 2021). However, these methods all rely on views of race and gender in which the ideal relationship between groups is one of sameness rather than difference. According to Drage and Mackereith (2022), addressing and solving inequities in resume screening mediated by AI or LLMs requires accounting for structural power imbalances that underpin the conceptualization and use of these tools.

The primary limitations of this study were data-related. The resumes may diverge from those being used by real-world applicants due to the need to truncate them for computational feasibility. Furthermore, we relied on external tools to determine occupation categories for each document. The codes may not be as accurate as manual coding, which limits the conclusions which can be drawn for particular occupations.

The results presented here exemplify the risks associated with using LLMs for resume screening. They consistently

replicate existing societal patterns of discrimination, further disadvantaging the groups already experiencing inequity in resume screening (Bertrand and Mullainathan 2004; Pager 2003). Additionally, screening outcomes are highly variable based on low-level features such as name frequency and resume length. Rather than LLMs having the potential to counteract people’s unconscious biases, we find these are reinforced in ways that are unpredictable and difficult to control in a real-world resume screening setting. Therefore, policy and mechanisms to comprehensively audit resume screening systems, whether proprietary or open source, are essential in order to evaluate their fairness and improve or remove these systems accordingly.

Future Work

Given the novelty of using LLMs as resume screeners, there is still much work to do in documenting their potential risks and improving their transparency in ways that can potentially help to identify and reduce bias and discrimination in hiring scenarios. Areas for future research include assessing additional MTEs, increasing the diversity of social group signals in resumes, as well as the range of social groups investigated. This study was limited to analyzing only two of the most commonly studied race (White and Black) and gender (male and female) groups via associated names. Hiring discrimination is not limited to these groups or signals, thus it will be important to investigate additional groups in order to fully quantify the risks in using LLMs for hiring. Finally, investigating realistic variations in resume length is an important direction as well.

Conclusion

We proposed using retrieval to simulate resume screening via LLMs to investigate the potential for biased outcomes. We found that the models do not exhibit occupational biases, but rather reinforce societal “defaults,” such as the preference for White and male identities. Further intersectional analyses showed that Black males are at the greatest disadvantage, and three hypotheses of intersectionality were also confirmed in this setting. Finally, we investigated how features like resume length and name frequency can also impact biased outcomes.

Ethical Considerations Statement

Resume screening using LLMs can be potentially difficult to research and audit ethically. Of primary importance is the preservation of privacy and confidentiality when using documents, such as resumes, which contain large amounts of identifiable information. Researchers interested in transparency and reproducibility must negotiate tensions between the distribution and use of documents which accurately reflect signals of identity such as race and gender as they would be present in real-world resumes, while also preserving privacy of the candidates represented by the documents.

Additionally, audit studies primarily represent race and gender through names, which is a reductive and incomplete way of representing these facets of identity. Despite its limitations, it is currently one of the main approaches to study

the impact of social groups and their protected attributes. As more sophisticated ways are developed to more accurately represent protected groups (including, but not limited to the race and gender groups investigated here), this study should be extended to investigate the impact of LLMs use as resume screeners on additional populations.

Acknowledgements

We are grateful to the anonymous reviewers for their helpful feedback. This work was supported by the U.S. National Institute of Standards and Technology (NIST) Grant 60NANB23D194. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of NIST.

References

- Anil, C.; Wu, Y.; Andreassen, A.; Lewkowycz, A.; Misra, V.; Ramasesh, V.; Slone, A.; Gur-Ari, G.; Dyer, E.; and Neyshabur, B. 2022. Exploring length generalization in large language models. *Advances in Neural Information Processing Systems*, 35: 38546–38556.
- Baert, S. 2018. *Hiring discrimination: An overview of (almost) all correspondence experiments since 2005*. Springer.
- Bertrand, M.; and Mullainathan, S. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review*, 94(4): 991–1013.
- Bogen, M.; and Rieke, A. 2018. Help wanted: An examination of hiring algorithms, equity, and bias.
- Caliskan, A.; Ajay, P. P.; Charlesworth, T.; Wolfe, R.; and Banaji, M. R. 2022. Gender Bias in Word Embeddings: A Comprehensive Analysis of Frequency, Syntax, and Semantics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '22, 156–170. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392471.
- Census, N. 2024. What are the 5,000 Most Common Last Names in the U.S.? — namecensus.com. <https://namecensus.com/last-names/>. [Accessed 28-04-2024].
- Chen, Z. 2023. Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications*, 10(1): 1–12.
- Cheng, M.; Durmus, E.; and Jurafsky, D. 2023. Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1504–1532. Toronto, Canada: Association for Computational Linguistics.
- Dastin, J. 2018. Insight - Amazon scraps secret AI recruiting tool that showed bias against women. <https://www.reuters.com/article/idUSKCN1MK0AG/>. [Accessed 28-04-2024].
- Deshpande, K. V.; Pan, S.; and Foulds, J. R. 2020. Mitigating Demographic Bias in AI-based Resume Filtering. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, UMAP '20 Adjunct, 268–275. New York, NY, USA: Association for Computing Machinery. ISBN 9781450379502.
- Drage, E.; and Mackereth, K. 2022. Does AI debias recruitment? Race, gender, and AI's "eradication of difference". *Philosophy & technology*, 35(4): 89.
- Elder, E. M.; and Hayes, M. 2023. Signaling race, ethnicity, and gender with names: Challenges and recommendations. *The Journal of Politics*, 85(2): 764–770.
- Esiobu, D.; Tan, X.; Hosseini, S.; Ung, M.; Zhang, Y.; Fernandes, J.; Dwivedi-Yu, J.; Presani, E.; Williams, A.; and Smith, E. 2023. ROBBIE: Robust Bias Evaluation of Large Generative Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 3764–3814. Singapore: Association for Computational Linguistics.
- Franke, T. M.; Ho, T.; and Christie, C. A. 2012. The chi-square test: Often used and more often misinterpreted. *American journal of evaluation*, 33(3): 448–458.
- Gerritse, E. J.; and de Vries, A. P. 2020. Effect of debiasing on information retrieval. In *Bias and Social Aspects in Search and Recommendation: First International Workshop, BIAS 2020, Lisbon, Portugal, April 14, Proceedings 1*, 35–42. Springer.
- Ghavami, N.; and Peplau, L. A. 2013. An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly*, 37(1): 113–127.
- Ghosh, S.; and Caliskan, A. 2023. 'Person' == Light-skinned, Western Man, and Sexualization of Women of Color: Stereotypes in Stable Diffusion. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 6971–6985. Singapore: Association for Computational Linguistics.
- Glazko, K.; Mohammed, Y.; Kosa, B.; Potluri, V.; and Mankoff, J. 2024. Identifying and Improving Disability Bias in GAI-Based Resume Screening. arXiv:2402.01732.
- González, M. J.; Cortina, C.; and Rodríguez, J. 2019. The role of gender stereotypes in hiring: A field experiment. *European Sociological Review*, 35(2): 187–204.
- Guo, W.; and Caliskan, A. 2021. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 122–133.
- Hipes, C.; Lucas, J.; Phelan, J. C.; and White, R. C. 2016. The stigma of mental illness in the labor market. *Social Science Research*, 56: 16–25.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. I.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jones, E.; and Steinhardt, J. 2022. Capturing failures of large language models via human cognitive biases. *Advances in Neural Information Processing Systems*, 35: 11785–11799.

- Kirk, H. R.; Jun, Y.; Volpin, F.; Iqbal, H.; Benussi, E.; Dreyer, F.; Shtedritski, A.; and Asano, Y. 2021. Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. *Advances in neural information processing systems*, 34: 2611–2624.
- Kotek, H.; Dockum, R.; and Sun, D. 2023. Gender bias and stereotypes in Large Language Models. In *Proceedings of The ACM Collective Intelligence Conference, CI '23*, 12–24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701139.
- Lahey, J. N. 2008. Age, women, and hiring: An experimental study. *Journal of Human resources*, 43(1): 30–56.
- Li, Y.; and Goel, S. 2024. Making It Possible for the Auditing of AI: A Systematic Review of AI Audits and AI Auditability. *Information Systems Frontiers*, 1–31.
- Liu, J.; Min, S.; Zettlemoyer, L.; Choi, Y.; and Hajishirzi, H. 2024. Infini-gram: Scaling Unbounded n-gram Language Models to a Trillion Tokens. *arXiv preprint arXiv:2401.17377*.
- Meng, R.; Liu, Y.; Rayhan Joty, S.; Xiong, C.; Zhou, Y.; and Yavuz, S. 2024. SFR-Embedding-Mistral: Enhance Text Retrieval with Transfer Learning. *Salesforce AI Research Blog*.
- Mishel, E. 2016. Discrimination against queer women in the US workforce: A résumé audit study. *Socius*, 2: 2378023115621316.
- Muennighoff, N.; Su, H.; Wang, L.; Yang, N.; Wei, F.; Yu, T.; Singh, A.; and Kiela, D. 2024. Generative representational instruction tuning. *arXiv preprint arXiv:2402.09906*.
- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2022. MTEB: Massive Text Embedding Benchmark. *arXiv preprint arXiv:2210.07316*.
- Narayanan Venkit, P. 2023. Towards a Holistic Approach: Understanding Sociodemographic Biases in NLP Models using an Interdisciplinary Lens. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, AIES '23*, 1004–1005. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.
- Neumark, D.; Burn, I.; and Button, P. 2016. Experimental age discrimination evidence and the Heckman critique. *American Economic Review*, 106(5): 303–308.
- Neumark, D.; Burn, I.; and Button, P. 2019. Is it harder for older workers to find jobs? New and improved evidence from a field experiment. *Journal of Political Economy*, 127(2): 922–970.
- Ouyang, S.; Wang, S.; Liu, Y.; Zhong, M.; Jiao, Y.; Iyer, D.; Pryzant, R.; Zhu, C.; Ji, H.; and Han, J. 2023. The Shifted and The Overlooked: A Task-oriented Investigation of User-GPT Interactions. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2375–2393. Singapore: Association for Computational Linguistics.
- Pager, D. 2003. The mark of a criminal record. *American journal of sociology*, 108(5): 937–975.
- Parasurama, P.; and Sedoc, J. 2021. Degendering Resumes for Fair Algorithmic Resume Screening. *arXiv preprint arXiv:2112.08910*.
- Parasurama, P.; and Sedoc, J. 2022. Gendered information in resumes and its role in algorithmic and human hiring bias. In *Academy of Management Proceedings*, volume 2022, 17133. Academy of Management Briarcliff Manor, NY 10510.
- Raghavan, M.; Barocas, S.; Kleinberg, J.; and Levy, K. 2020. Mitigating bias in algorithmic hiring: evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, 469–481. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Rhea, A.; Markey, K.; D’Arinzo, L.; Schellmann, H.; Sloane, M.; Squires, P.; and Stoyanovich, J. 2022. Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring: Results of an Audit. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, AIES '22*, 572–587. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392471.
- Sánchez-Monedero, J.; Dencik, L.; and Edwards, L. 2020. What does it mean to ‘solve’ the problem of discrimination in hiring? social, technical and legal perspectives from the UK on automated hiring systems. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, 458–468. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Schellmann, H. 2024. *The algorithm*. Hachette Books.
- Schmitz, M.; Forst, L.; et al. 2016. Industry and occupation in the electronic health record: an investigation of the National Institute for Occupational Safety and Health Industry and Occupation Computerized Coding System. *JMIR medical informatics*, 4(1): e4839.
- Sciar, M.; Choi, Y.; Tsvetkov, Y.; and Suhr, A. 2023. Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. *arXiv preprint arXiv:2310.11324*.
- Soldaini, L.; Kinney, R.; Bhagia, A.; Schwenk, D.; Atkinson, D.; Authur, R.; Bogin, B.; Chandu, K.; Dumas, J.; Elazar, Y.; Hofmann, V.; Jha, A. H.; Kumar, S.; Lucy, L.; Lyu, X.; Lambert, N.; Magnusson, I.; Morrison, J.; Muennighoff, N.; Naik, A.; Nam, C.; Peters, M. E.; Ravichander, A.; Richardson, K.; Shen, Z.; Strubell, E.; Subramani, N.; Tafjord, O.; Walsh, P.; Zettlemoyer, L.; Smith, N. A.; Hajishirzi, H.; Beltagy, I.; Groeneveld, D.; Dodge, J.; and Lo, K. 2024. Dolma: an Open Corpus of Three Trillion Tokens for Language Model Pretraining Research. *arXiv preprint*.
- Sundararaman, D.; and Subramanian, V. 2022. Debiasing Gender Bias in Information Retrieval Models. *arXiv preprint arXiv:2208.01755*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2023. Improving text embeddings with large language models. *arXiv preprint arXiv:2401.00368*.

Wilson, C.; Ghosh, A.; Jiang, S.; Mislove, A.; Baker, L.; Szary, J.; Trindel, K.; and Polli, F. 2021. Building and Auditing Fair Algorithms: A Case Study in Candidate Screening. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, 666–677. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.

Wolfe, R.; and Caliskan, A. 2021. Low frequency names exhibit bias and overfitting in contextualizing language models. *arXiv preprint arXiv:2110.00672*.

Wolfe, R.; and Caliskan, A. 2022. American == White in Multimodal Language-and-Image AI. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, AIES '22*, 800–812. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392471.

Wright, B. R.; Wallace, M.; Bailey, J.; and Hyde, A. 2013. Religious affiliation and hiring discrimination in New England: A field experiment. *Research in Social Stratification and Mobility*, 34: 111–126.

Yin, L.; Alba, D.; and Nicoletti, L. 2024. OpenAI’s GPT Is a Recruiter’s Dream Tool. Tests Show There’s Racial Bias.

Zhao, W. X.; Liu, J.; Ren, R.; and Wen, J.-R. 2024. Dense Text Retrieval Based on Pretrained Language Models: A Survey. *ACM Trans. Inf. Syst.*, 42(4).